arXiv:1702.04307v1 [cs.DS] 14 Feb 2017

Approximating the Held-Karp Bound for Metric TSP in Nearly-Linear Time*

Chandra Chekuri Kent Quanrud

February 15, 2017

Abstract

We give a nearly linear-time randomized approximation scheme for the Held-Karp bound [Held and Karp, 1970] for Metric-TSP. Formally, give an undirected edge-weighted graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ on m edges and $\epsilon > 0$, the algorithm outputs in $O(m \operatorname{polylog}(m)/\epsilon^2)$ time, with high probability, a $(1+\epsilon)$-approximation to the Held-Karp bound on the Metric-TSP instance induced by the shortest path metric on $\mathcal{G}$. We substantially improve upon the $O(m^2 \log^2(m)/\epsilon^2)$ run time achieved previously by Garg and Khandekar. The algorithm builds on our recent framework for accelerating multiplicative weight update based methods for implicit fractional packing problems.

*Department of Computer Science, University of Illinois, Urbana, IL 61820. {chekuri,quanrud2}@illinois.edu.
Work on this paper partially supported by NSF grant CCF-1526799.

1 Introduction

The *Traveling Salesman Problem* (TSP) is a central problem in discrete and combinatorial optimization, and a wellspring of inspiration for ideas in optimization, mathematical programming and theoretical computer science. Cook's recent book [2014] gives an introduction to the problem, its history, and general appeal. See also Gutin and Punnen [2006], Applegate, Bixby, Chvatal, and Cook [2011], and Lawler, Lenstra, Rinnooy-Kan, and Shmoys [1985] for book-length treatments of TSP and its variants.

Formally, the input to TSP is a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ equipped with positive edge costs $c : \mathcal{E} \to \mathbb{R}_{>0}$. The goal is to find a minimum cost Hamiltonian cycle in $\mathcal{G}$. In this paper we focus on TSP in undirected graphs. Checking whether a given graph has a Hamiltonian cycle is a classical NP-Complete decision problem, and hence TSP is not only NP-Hard but also inapproximable. For this theoretical reason, as well as many practical applications, a special case of TSP called Metric-TSP is extensively studied. In Metric-TSP, $\mathcal{G}$ is a complete graph K_n and c obeys the triangle inequality $c_{uv} \leq c_{uw} + c_{wv}$ for all $u, v, w \in \mathcal{V}$. An alternative interpretation of Metric-TSP is that the input is an edge-weighted graph and the goal is to find a minimum-cost *tour* of $\mathcal{G}$. A tour is a closed walk that visits all the vertices and returns to the starting point. In other words, Metric-TSP is the relaxation of TSP where a vertex is allowed to be visited more than once. The graph-based view of Metric-TSP allows one to specify the metric on $\mathcal{V}$ implicitly and *sparsely*.

Unlike TSP, which is inapproximable, Metric-TSP admits a constant factor approximation. The classical algorithm of Christofides [1976] yields a 3/2-approximation. On the other hand it is known that Metric-TSP is APX-Hard and hence does not admit a PTAS (Lampis [2012] showed that there is no $\frac{185}{184}$-approximation unless $P = NP$). An outstanding open problem in approximation algorithms is to improve the bound of 3/2. A well-known conjecture states that the worst-case integrality gap of the *Subtour-Elimination LP* formulated by Dantzig, Fulkerson, and Johnson [1954] is 4/3 (see [Goemans, 1995]). There has been exciting recent progress on this conjecture and several related problems; we refer the reader to an excellent survey by Vygen [2012]. The Subtour Elimination LP for TSP is described below and is based on variables y_e for each edge $e \in \mathcal{E}$. In the following, let $\mathcal{C}(S)$ (resp. $\mathcal{C}(v)$) denotes the set of edges crossing the set of vertices $U \subseteq \mathcal{V}$ (resp. the vertex $v \in \mathcal{V}$).

$$\begin{aligned} \mathrm{SE}(\mathcal{G}, c) = \min \quad & \langle c, y \rangle & \\ \text{s.t.} \quad & \textstyle\sum_{e \in \mathcal{C}(v)} y_e = 2 & \text{for all } v \in \mathcal{V}, \\ & \textstyle\sum_{e \in \mathcal{C}(U)} y_e \geq 2 & \text{for all } \emptyset \subsetneq U \subsetneq \mathcal{V}, \\ \text{and} \quad & y_e \in [0, 1] & \text{for all } e \in \mathcal{E}. \end{aligned}$$

The first set of constraints consists of degree constraints that require each vertex to be incident to two edges (in the integral setting). The second set of constraints force connectivity, hence the name "subtour elimination". The LP provides a lower bound for TSP, and in order to apply it to an instance of Metric-TSP defined by $\mathcal{G}$, one needs to use it on the metric completion of $\mathcal{G}$.

A problem closely related to Metric-TSP is the 2-*edge-connected spanning subgraph problem* (2ECSS). In 2ECSS the input is an edge-weighted graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, and the goal is to find a minimum cost subgraph of $\mathcal{G}$ that is 2-edge-connected. We focus on the simpler version where an edge is allowed to be used more than once. A natural LP relaxation for 2ECSS is described below on the left. We have a variable y_e for each edge $e \in \mathcal{E}$, and constraints which ensure that each cut has at least two edges crossing it. We also describe the dual LP on the right which corresponds to a maximum packing of cuts into the edge costs. In the following, let $\mathcal{C} \subseteq 2^{\mathcal{E}}$ denote the family of all cuts in $\mathcal{G}$. (For technical reasons, we prefer to treat cuts as sets of edges.)

$$\begin{aligned} &\mathsf{2ECSS}(\mathcal{G}, c) = \min\langle c, y\rangle \\ &\quad \text{s.t.} \sum_{e \in C} y_e \geq 2 \qquad \forall C \in \mathcal{C} \\ &\quad \text{and } y_e \geq 0 \qquad \forall e \in \mathcal{E} \end{aligned}$$

$$\begin{aligned} &\mathsf{2ECSSD}(\mathcal{G}, c) = \max\ 2\langle \mathbb{1}, x\rangle \\ &\quad \text{s.t.} \sum_{C \ni e} x_C \leq 2 \qquad \forall C \in \mathcal{C} \\ &\quad \text{and } x_C \geq 0 \qquad \forall C \in \mathcal{C} \end{aligned}$$

Cunningham (see [Monma, Munson, and Pulleyblank, 1990]) and Goemans and Bertsimas [1993] observed that for any edge-weighted graph $(\mathcal{G}, c)$, the optimum value of the Subtour Elimination LP for the metric completion of $(\mathcal{G}, c)$ coincides with the 2ECSS LP for $(\mathcal{G}, c)$. The advantage of this connection is twofold. First, the 2ECSS relaxation is a pure covering LP, and its dual is a pure packing LP. Second, the 2ECSS formulation works with the underlying graph $(\mathcal{G}, c)$ directly instead of the metric completion.

Although the LP relaxations provide good lower bounds for Metric-TSP, they require the Ellipsoid method to solve in polynomial-time. One can write polynomial-sized extended formulations using flow variables and the maxflow-mincut theorem, but the number of variables becomes quite large. Held and Karp [1970] provided an alternative lower bound for TSP via the notion of *one-trees*. Interestingly, they showed that their lower bound coincides with $\mathrm{SE}(\mathcal{G}, c)$. The advantage of the Held-Karp bound is that it can computed via a simple iterative procedure that relies on minimum spanning tree algorithms. However, there does not appear to be a polynomial-time implementation with good guarantees on the quality of the approximation.

In the rest of the paper we focus on Metric-TSP. For the sake of brevity, we refer to the Held-Karp bound for the metric completion of $(\mathcal{G}, c)$ as simply the Held-Karp bound for $(\mathcal{G}, c)$. How fast can the Held-Karp bound for a given instance be computed? A fast algorithm has several applications. The LP solution is a starting point for some candidate approximation algorithms to prove the 4/3 conjecture. The lower bound is also an important ingredient in integer programming techniques such as branch and bound. Plotkin, Shmoys, and Tardos [1995], in their influential paper on fast approximation schemes for packing and covering LPs via Lagrangian relaxation methods, showed that a $(1+\epsilon)$-approximation for the Held-Karp bound for Metric-TSP can be computed in $O\big(n^4 \log^6 n/\epsilon^2\big)$ randomized time. They relied on an algorithm for computing the global minimum cut[1]. Subsequently, Garg and Khandekar obtained a $(1+\epsilon)$-approximation in $O(m^2 \log^2 m/\epsilon^2)$ time and they relied on algorithms for minimum-cost branchings (see [Khandekar, 2004]).

Our main result, stated below, obtains a near-linear running time for a $(1+\epsilon)$-approximation, substantially improving the previous bounds.

Theorem 1.1. *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be an undirected graph with $|\mathcal{E}| = m$ edges and $|\mathcal{V}| = n$ vertices, and positive edge weights $c : \mathcal{E} \to \mathbb{R}_{>0}$. For any fixed $\epsilon > 0$, there exists a randomized algorithm that computes a $(1+\epsilon)$-approximation to the Held-Karp lower bound for the* Metric-TSP *instance on $(\mathcal{G}, c)$ in $O(m \log^4 n/\epsilon^2)$ time. The algorithm succeeds with high probability.*

For fixed ϵ, the running time is asymptotically faster than the time to compute or even write down the metric completion of $(\mathcal{G}, c)$. In typical approximation algorithms that rely on mathematical programming relaxations, the bottleneck for the running time is solving the relaxation. Surprisingly, for algorithms solving Metric-TSP via the Held-Karp bound, the bottleneck is no longer solving the relaxation (albeit we only find a $(1+\epsilon)$-approximation and do not guarantee a basic feasible solution).

[1]Their scheme can in fact be implemented in randomized $O(m^2 \log^4 m/\epsilon^2)$ time using subsequent developments in minimum cut algorithms and width reduction techniques.

We remark that our algorithm can also be used for geometric instances. Given a Metric-TSP instance induced by a set P of n points in $\mathbb{R}^d$ equipped with the Euclidean metric, and an $\epsilon > 0$, one can compute in $O\big(n \log n + n\epsilon^{-d}\big)$ time, a graph on P that is a $(1+\epsilon)$-approximate spanner, and has $O\big(n\epsilon^{-d}\big)$ edges [Har-Peled, 2011, Theorem 3.12]; this also extends to points with doubling dimension d. Thus, the Held-Karp bound for P can be computed in $\tilde{O}\big(n\epsilon^{-(d+2)}\big)$ time. It may be feasible to improve the dependence on ϵ since we do not need the full strength of a $(1+\epsilon)$-spanner. A $(1+\epsilon)$-approximation for Metric-TSP (in terms of the actual integral optimum) in Euclidean spaces can be computed in $O(2^{(d/\epsilon)^{O(d)}} n)$ time [Bartal and Gottlieb, 2013, Rao and Smith, 1998]. The exponential dependence on $1/\epsilon$ is necessary (unless $P = NP$).

Implicit in the proof of Theorem 1.1 is a data structure that maintains a $(1+\epsilon)$-approximate global mincut of a partially dynamic graph where edge-weights only increase. We achieve amortized polylogarithmic time for edge weight updates and maintain a compact representation of an approximately minimum min-cut. The actual edges of the approximate minimum cut can be reported in polylogarithmic time plus constant time per edge reported. This is not the main focus of the paper and we do not explicitly state or work out a formal interface to this data structure.

Integrated design of the algorithm. Our algorithm is based on the multiplicative weight update framework (MWU), and like Plotkin, Shmoys, and Tardos [1995], we approximate the pure packing LP 2ECSSD. Each iteration requires an oracle for computing the global minimum cut in an undirected graph. A single minimum cut computation takes randomized near-linear-time via the algorithm of [Karger, 2000], and the MWU framework requires $\tilde{\Omega}(m/\epsilon^2)$ iterations. Suprisingly, the whole algorithm can be implemented to run in roughly the same time as that required to compute one global mincut.

While the full algorithm is fairly involved, the high-level design is directed by relatively basic ideas developed in recent work by the authors [Chekuri and Quanrud, 2017]. Inspired by earlier work of Mądry [2010] and Young [2014], we accelerate MWU-based algorithms for some implicit packing problems with the careful interplay of two data structures. The first data structure maintains a minimum cost object of interest (here the global minimum cut) by (partially) dynamic techniques, rather than recompute the object from scratch in every iteration. The second data structure applies the multiplicative weight update in a lazy fashion that can be amortized efficiently against the weights of the constraints. The two data structures need to be appropriately meshed to obtain faster running times, and this meshing depends very much on the problem specifics.

While we do inherit some basic ideas and techniques from this framework, the problem here is more sophisticated than those considered in [Chekuri and Quanrud, 2017]. The first component in this paper is a fast dynamic data structure for maintaining a $(1+\epsilon)$-approximate global minimum cut of a weighted graph whose edge weights are only increasing. We achieve an amortized polylogarithmic update time by a careful adaptation of the randomized near-linear-time minimum cut algorithm of Karger [2000] that relies on approximate tree packings. This data structure is developed with careful consideration of the MWU framework; for example, we only need to compute an approximate tree packing $\tilde{O}\big(1/\epsilon^2\big)$ times (rather than in every iteration) because of the monotonicity of the weight updates and standard upper bounds on the total growth of the edge weights.

The second technical ingredient is a data structure for applying a multiplicative weight update to each edge in the approximately minimum cuts selected in each iteration. The basic difficulty here is that we cannot afford to touch each edge in the cut. While this task suggests a lazy weight update similar to [Chekuri and Quanrud, 2017], these techniques require a compact representation of the minimum cuts. It appears difficult to develop in isolation a data structure that can apply multiplicative weight updates along any (approximately) minimum cut. However, additional nice properties of the cuts generated by the first dynamic data structure enable a clean interaction

with lazy weight updates. We develop an efficient data structure for weight updates that is tightly coupled with the data structure generating the approximately minimum cuts.

There are many technical details that the above overview necessarily skips. We develop the overall algorithm in a methodical and modular fashion in the rest of the paper.

1.1 Related Work

Our main result has connections to several important and recent directions in algorithms research. We point out the high-level aspects and leave a more detailed discussion to a later version of the paper.

In one direction there has been a surge of interest in faster (approximate) algorithms for classical problems in combinatorial optimization including flows, cuts, matchings and linear programming to name a few. These are motivated by not only theoretical considerations and techniques but also practical applications with very large data sizes. Many of these are based on the interplay between discrete and continuous techniques and there have been several breakthroughs, too many to list here. One prominent recent example of a problem that admits a near-linear-time approximation scheme is maximum flow in undirected graphs [Peng, 2016] building on many previous tools and techniques. Our work adds to the list of basic problems that admit a near-linear-time approximation scheme.

In another direction there has been exciting work on approximation algorithms for TSP and its variants including Metric-TSP, Graphic-TSP, Asymmetric-TSP, TSP-Path to name a few. Several new algorithms, techniques, connections and problems have arisen out of this work. Instead of recapping the many results we refer the reader to the survey [Vygen, 2012]. Our result adds to this literature and suggests that one can obtain not only improved approximation but also much faster approximations than what was believed possible. We mention that several of the algorithms based on the LP relaxation decompose the LP solution x into a convex combination of spanning trees (after scaling x down appropriately), randomly sample one of them, and then complete the tree to a tour ala Chritofides heuristic. Our recent work [Chekuri and Quanrud, 2017] shows how to decompose x in the spanning tree polytope into a $(1-\epsilon)$-approximate convex combination of spanning trees in near-linear time. It may be feasible to combine additional ingredients, such as a fast algorithm to compute an approximate minimum cost matching, to implement the whole algorithm in near-linear time.

In terms of techniques, our work falls within the broad framework of improving iterative algorithms via data structures. Our recent work showed that MWU based algorithms, especially for implicit packing and covering problems, have the potential to be substantially improved by taking advantage of data structures. The key, as we described already, is to combine multiple data structures together while exploiting the flexibility of the MWU framework. Although not a particularly novel idea (see the work of Mądry [2010] for applications to multicommodity flow and the work of Agarwal and Pan [2014] for geometric covering problems), our work demonstrates concrete and interesting problems to highlight the sufficient conditions under which this is possible. Lagrangian-relaxation based approximation schemes for solving special classes of linear programs have been studied for several decades with a systematic investigation started in Plotkin, Shmoys, and Tardos [1995] and Grigoriadis and Khachiyan [1994] following earlier work on multicommodity flows. Since then there have been many ideas, refinements and applications. We refer the reader to [Koufogiannakis and Young, 2014], [Young, 2014], [Chekuri, Jayram, and Vondrák, 2015] for some recent papers on MWU based algorithms which also have additional pointers, and to the survey [Arora, Hazan, and Kale, 2012] for the broader applicability of MWU in theoretical computer science. Accelerated gradient descent methods have recently resulted in fast randomized near-linear-time algorithms for

explicit fractional packing and covering LPs; these algorithms improved the dependence of the running time on ϵ to $1/\epsilon$ from $1/\epsilon^2$. See [Allen-Zhu and Orecchia, 2015] and [Wang, Rao, and Mahoney, 2015].

Our work here builds extensively on Karger's randomized nearly-linear-time mincut algorithm [2000]. As mentioned already, we adapt his algorithm to the partially dynamic setting. [Thorup, 2007] also builds on Karger's tree packing ideas to develop a *fully* dynamic mincut algorithm. Thorup's algorithm is rather involved and is slower than what we are able to achieve for the partial dynamic setting; he achieves an update time of $\widetilde{O}(\sqrt{n})$ while we achieve polylogarithmic time. There are other obstacles to integrating his ideas and data structure to the needs of the MWU framework: for instance, we need to consider all $(1+\epsilon)$-approximate mincuts, and we also need to interface cleanly with the needs of the data structure to maintain weights lazily.

Some of the data structures in this paper rely on Euler tour representations of spanning trees. The Euler tour representation imposes a particular linear order on the vertices by which the edges of the underlying graph can be interpreted as intervals. The Euler tour representation was introduced by Tarjan and Vishkin [1985] and has seen applications in other dynamic data structures, such as the work by Henzinger and King [1999] among others.

Organization. Section 2 gives a high-level overview of a first MWU-based algorithm, highlighting some important properties of the MWU framework that arise in more sophisticated arguments later. Section 3 summarizes the key properties of tree packings that underlie Karger's mincut algorithm, and explains their use in our MWU algorithm via the notion of epochs. Section 4 describes an efficient subroutine, via appropriate data structures, to find all approximate mincuts (in the partial dynamic setting) induced by the spanning trees of a tree packing. Section 5 describes the data structure that implement the weights of the edges in a lazy fashion, building on Euler tours of spanning trees, range trees, and some of our prior work. In Section 6, we outline a formal proof of Theorem 1.1.

2 MWU based Algorithm and Overview

Given an edge-weighted graph $(\mathcal{G} = (\mathcal{V}, \mathcal{E}), c)$ our goal is to find a $(1-\epsilon)$-approximation to the optimal value of the LP 2ECSSD. We later describe how to obtain an approximate solution to the primal LP 2ECSS. Note that the LP is a packing LP of the form $\max\langle \mathbb{1}, x\rangle : Ax \leq c$ with an exponential number of variables corresponding to the cuts, but only m non-trivial constraints corresponding to the edges. The MWU framework can be used to obtain a $(1-\epsilon)$-approximation for fractional packing. Here we pack cuts into the edge capacities c. The MWU framework reduces packing cuts to finding global minimum cuts as follows. The framework starts with an empty solution $x = \mathbb{0}^{\mathcal{C}}$ and maintains edge weights $w : \mathcal{E} \to \mathbb{R}_{\geq 0}$, initialized to $1/c$ and *non-decreasing* over the course of the algorithm. Each iteration, the framework finds the minimum cut $C = \arg\min_{C' \in \mathcal{C}} \overline{w}(C')$ w/r/t w, where $\overline{w}(C') \stackrel{\text{def}}{=} \sum_{e \in C'} w_e$ denotes the total weight of a cut. The framework adds a fraction of C to the solution x, and updates the weight of every edge in the cut in a multiplicative fashion such that the weight of an edge is exponential in its load $x(e)/c(e)$. The weight updates steer the algorithm away from reusing highly loaded edges. The MWU framework guarantees that the output x will have objective value $\langle \mathbb{1}, x\rangle \geq \text{OPT}$ while satisfying $\sum_{C \ni e} x_C \leq (1+O(\epsilon))c_e$ for every edge e. Scaling down x by a $(1+O(\epsilon))$-factor gives a $(1-O(\epsilon))$-approximation satisfying all packing constraints.

An actual implementation of the above high-level idea needs to specify the step size in each iteration and the precise weight-update. The non-uniform increments idea of Garg and Könemann [2007] gives rise to a width-independent bound on the number of iterations, namely $O(m \log m/\epsilon^2)$

```
held-karp-1(G = (V, E), c, ε)
  x ← 0^C,  w ← 1/c,  t ← 0,  η ← ε/ln m
  while t < 1
    C ← arg min_{C∈C} w̄(C)  // compute the minimum cut
    y ← (⟨w, c⟩ / w̄(C)) · 1_C  // y = arg max { ⟨1, z⟩ : z ∈ R^C_{≥0} and Σ_{e∈E} w_e Σ_{C∋e} z_C ≤ ⟨w, c⟩ }
    γ ← min_{e∈C} c_e
    δ ← εγ/η  // such that δ Σ_{C∋e} y_C ≤ (ε/η) c_e for all e ∈ E
    x ← x + δy  // x is maintained implicitly
    for all e ∈ C
      w_e ← w_e exp(δη/c_e) = w_e exp(εγ/c_e)  // increase the weight of edge e
    t ← t + δ
  end while
  return x
```

Figure 1: An exact (and slow) implementation of the MWU framework applied to packing cuts.

---

in our setting. Our algorithms follow the specific notation and scheme of Chekuri et al. [2015], which tracks the algorithm's progress by a "time" variable t from 0 to 1 increasing in non-uniform steps. A step of size δ in an iteration corresponds to adding $\delta\beta 1_C$ to the current solution x where 1_C is the characteristic vector of the mincut C found in the iteration, and $\beta = \langle w, c\rangle / \overline{w}(C)$ greedily takes as much of the cut as can fit in the budget $\langle w, c\rangle$. The process is controlled by a parameter η, which when set to $O(\ln m/\epsilon)$ guarantees both the width-independent running time bound as well as the $(1-\epsilon)$-approximation.

Both the correctness and the number of iterations are derived by a careful analysis of the edge weights w. We review the argument that bounds the number of iterations. At the beginning of the algorithm, when every edge e has weight $w_e = 1/c_e$, we have $\langle w, c\rangle = m$. The weights monotonically increase, and standard proofs show that this inner product is bounded above by $\langle w, c\rangle \leq \tilde{O}\big(m^{O(1/\epsilon)}\big)$. The edge weight increases are carefully calibrated so that at least one edge increases by a $(1+\epsilon)$-multiplicative factor in each iteration. As the upper bound on $\langle w, c\rangle$ is an upper bound on $w_e c_e$ for any edge e, and the initial weight of an edge e is $1/c_e$, an edge weight w_e can increase by a $(1+\epsilon)$-factor at most $\log_{(1+\epsilon)}\big(m^{O(1/\epsilon)}\big) = O\big(\log(m)/\epsilon^2\big)$ times. Charging each iteration to an edge weight increased by a $(1+\epsilon)$-factor, the algorithm terminates in $\tilde{O}\big(m/\epsilon^2\big)$ iterations.

A direct implementation of the MWU framework, `held-karp-1`, is given in Figure 1. In addition to the input $(\mathcal{G}, c, \epsilon)$, there is an additional parameter $\eta > 0$ that dampens the step size at each iteration. Standard analysis shows that the appropriate choice of η is about $\ln(m)/\epsilon$. Each iteration requires the minimum global cut w/r/t the edge weights w, which can be found in $\tilde{O}(m)$ time (with high probability) by the randomized algorithm of Karger [2000]. Calling Karger's algorithm in each iteration gives a quadratic running time of $\tilde{O}\big(m^2/\epsilon^2\big)$.

Approximation in the framework. The MWU framework is robust to approximation in several ways, especially in the setting of packing and covering where non-negativity is quite helpful. For example, it suffices to find a $(1+O(\epsilon))$-multiplicative approximation of the minimum cut, maintain the edge weights to a $(1+O(\epsilon))$-multiplicative approximation of their true edge weights, and so

held-karp-2($\mathcal{G} = (\mathcal{V}, \mathcal{E}), c, \epsilon$)

$x \leftarrow \mathbb{0}^{\mathcal{C}}$, $w \leftarrow 1/c$, $t \leftarrow 0$, $\eta \leftarrow \epsilon / \ln m$

$\lambda \leftarrow$ size of minimum cut w/r/t edge weights w

while $t < 1$

  while (a) $t < 1$ and

    (b) there is a cut $C \in \mathcal{C}$ s.t. $\overline{w}(c) \leq (1+\epsilon)\lambda$

    $y \leftarrow \dfrac{\langle w, c \rangle}{\overline{w}(C)} \cdot 1_C$ // $\langle \mathbb{1}, y \rangle \leq (1+\epsilon) \max \left\{ \langle \mathbb{1}, z \rangle : z \in \mathbb{R}^{\mathcal{C}}_{\geq 0} \text{ and } \sum_{e \in E} w_e \sum_{C \ni e} z_C \leq \langle w, c \rangle \right\}$

    $\gamma \leftarrow \min_{e \in C} c_e$

    $\delta \leftarrow \dfrac{\epsilon \gamma}{\eta}$ // such that $\delta \sum_{C \ni e} y_C \leq \dfrac{\epsilon}{\eta} c_e$ for all $e \in \mathcal{E}$

    $x \leftarrow x + \delta y$

    for all $e \in C$

      $w_e \leftarrow \exp(\delta \eta / c_e) \cdot w_e = \exp(\epsilon \gamma / c_e) \cdot w_e$

    $t \leftarrow t + \delta$

  end while

  $\lambda \leftarrow (1+\epsilon)\lambda$ // start new epoch

end while

return x

Figure 2: A second implementation of the MWU framework applied to packing cuts that divides the iterations into epochs and seeks only approximately minimum cuts at each iteration.

---

forth, while retaining a $(1 + O(\epsilon))$-multiplicative approximation overall. We exploit this slack in two concrete ways that we describe at a high-level below.

Approximate mincuts. First, we only look for cuts that are within a $(1+O(\epsilon))$-approximate factor of the true minimum cut. To this end, we maintain a target cut value $\lambda > 0$, and maintain the invariant that there are no cuts of value strictly less than λ. We then look for cuts of value $\leq (1 + O(\epsilon))\lambda$ until we are sure that there are no more cuts of value $\leq (1+\epsilon)\lambda$. When we have certified that there are no cuts of value $\leq (1+\epsilon)\lambda$, we increase λ to $(1+\epsilon)\lambda$, and repeat. Each stretch of iterations with the same target value λ is called an *epoch*. Epochs were used by [Fleischer, 2000] for approximating fractional multicommodity flow. We show that *all* the approximately minimum cuts in a single epoch can be processed in $\tilde{O}(m)$ total time (plus some amortized work bounded by other techniques). If κ is the weight of the initial minimum cut (when $w = 1/c$), then the minimum cut size increases monotonically from κ to $\tilde{O}\big(m^{1/\epsilon}\kappa\big)$. If the first target cut value is $\lambda = \kappa$, then there are only $O(\ln m/\epsilon^2)$ epochs over the course of the entire algorithm.

Lazy weight updates. Even if we can concisely identify the approximate minimum cuts quickly, there is still the matter of updating the weights of all the edges in the cut quickly. A cut can have $\Omega(m)$ edges, and updating m edge weights in each of $\tilde{O}\big(m/\epsilon^2\big)$ iterations leads to a $\tilde{O}\big(m^2/\epsilon^2\big)$ running time. Instead of visiting each edge of the cut, we have a subroutine that simulates the weight increase to all the edges in the cut, but only does work for each edge whose (true) weight increases by a full $(1+\epsilon)$-multiplicative power. That is, for each edge e, we do work proportional (up to log factors) to the number of times w_e increases to the next integer power of $(1+\epsilon)$. By the previous discussion on weights, the number of such large updates is at most $O(\log m/\epsilon^2)$ per edge. This amortized efficiency comes at the expense of approximating the (true) weight of each

edge to a $(1 + O(\epsilon))$-multiplicative factor. Happily, an approximate minimum cut on approximate edge weights is still an approximate minimum cut to the true weights, so these approximations are tolerable. In [Chekuri and Quanrud, 2017] we isolated a specific lazy-weight scheme from Young [2014] into a data structure with a clean interface that we build upon here.

Obtaining a primal solution. Our MWU algorithm computes a $(1-\epsilon)$-approximate solution to 2ECSSD. Recall that the algorithm maintains a weight $w(e)$ for each edge e. Standard arguments show that one can recover from the evolving weights a $(1+\epsilon)$-approximate solution to the LP 2ECSS (which is the dual of 2ECCSD). A sketch is provided in the appendix.

3 Tree packings and epochs

Karger [2000] gave a randomized algorithm that finds the global minimum cut in a weighted and undirected graph with high probability in $\tilde{O}(m)$ time. Treating Karger's algorithm as a black box in every iteration leads to a quadratic running time of $\tilde{O}\big(m^2/\epsilon^2\big)$, so we open it up and adjust the techniques to our setting.

Tree packings. Karger [2000] departs from previous algorithms for minimum cut with an approach based on packing spanning trees. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be an undirected graph with positive edge weights $w : \mathcal{E} \to \mathbb{R}_{>0}$, and let $\mathcal{T} \subseteq 2^{\mathcal{E}}$ denote the family of spanning trees in $\mathcal{G}$. A *tree packing* is a nonnegatively weighted collection of spanning trees, $p : \mathcal{T} \to \mathbb{R}_{\geq 0}$, such that for any edge e, the total weight of trees containing e is at most c_e (i.e., $\sum_{T \ni e} p(T) \leq c_e$). Classical work of Tutte [1961] and Nash-Williams [1961] gives an exact characterization for the value of a maximum tree packing in a graph, which as an easy corollary implies it is at least half of the value of a mincut. If $C \in \mathcal{C}$ is a minimum cut and p is a maximum packing, then a tree $T \in \mathcal{T}$ selected randomly in proportion to its weight in the packing will share ≤ 2 edges with C in expectation. By Markov's inequality, T has strictly less than 3 edges in C with constant probability.

For a fixed spanning tree T, a *one-cut* in G induced by an edge $e \in T$ is the cut $\mathcal{C}(X)$ where X is the vertex set of one of the components of $T - e$. A *two-cut* induced by two edges $e_1, e_2 \in T$ is the following. Let X, Y, Z be the vertex sets of the three components of $T - \{e_1, e_2\}$ where X is only incident to e_1 and Z is only incident to e_2 and Y is incident to both e_1, e_2. Then the two-cut induced by e_1, e_2 is $\mathcal{C}(Y) = \mathcal{C}(X \cup Z)$. Thus, if T is a spanning tree sampled from a maximum tree packing, and C is a minimum cut, then C is either a one-cut or a two-cut with constant probability.

The probabilistic argument extends immediately to approximations. Let $\zeta \in (0, 1)$ be sufficiently small. If C is an approximately minimum cut with $\overline{w}(C) \leq (1+\zeta)\kappa$, and p is a tree packing of total weight $\geq (1-\zeta)\kappa/2$, then a random tree T sampled from p has $2+O(\zeta)$ edges from C in expectation, and strictly less than three edges in T with constant probability. Sampling $O(\log(1/\delta))$ trees from a tree packing amplifies the probability of a tree containing ≤ 2 edges of C from a constant to $\geq 1-\delta$. For constant $\zeta > 0$, a $(1-\zeta)$-approximate tree-packing can be computed in $\tilde{O}(m)$ time, either by applying $\tilde{O}(\kappa m)$-time tree packing algorithms [Gabow, 1995, Plotkin et al., 1995] to a randomly sparsified graph [Karger, 1998], or directly and deterministically in $\tilde{O}(m)$ time by a recent algorithm of Chekuri and Quanrud [2017].

Another consequence of Karger [2000] is that for $\zeta < 1/2$, the number of $(1+\zeta)$-approximate minimum cuts is at most n^2. By the union bound, if we select $O(\log n)$ trees at random from an approximately maximum tree packing, then with high probability *every* $(1+\zeta)$-approximate minimum cut is induced by one or two edges in one of the selected trees. In summary, we have the following.

Theorem 3.1 (Karger, 2000). *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be an undirected graph with edge capacities $c : \mathcal{E} \to \mathbb{R}_{>0}$, let $\delta \in (0,1)$, and let $\epsilon \in (0, 1/2)$. One can generate, in $O\big(m + n \log^3 n\big)$ time, $h = O(\log(n/\delta))$ spanning trees $T_1, T_2, \ldots, T_h$ such that with probability $\geq 1 - \delta$,* every *$(1+\epsilon)$-approximate minimum cut C has $|C \cap T_i| \leq 2$ for some tree T_i.*

Karger [2000] finds the minimum cut by checking, for each tree T_i, the minimum cut in G obtained by removing one or two edges from T_i in $\tilde{O}(m)$ time. (The details of this subroutine are reviewed in the following section.) Since there are $O(\log n)$ trees, this amounts to a near-linear-time algorithm. Whereas Karger's algorithm finds one minimum cut, we need to find many minimum cuts. Moreover, each time we find one minimum cut, the edge weights on that cut increase and so the underlying graph changes. We are faced with the challenge of outputting $\tilde{O}\big(m/\epsilon^2\big)$ cuts from a dynamically changing graph that in particular adapts to each cut output, all in roughly the same amount of time as a single execution of Karger's algorithm.

Epochs. We divide the problem into $O(\log m/\epsilon^2)$ epochs. Recall that at the start of an epoch we have a target value $\lambda > 0$ and are guaranteed that there are no cuts with value strictly less than λ. Our goal is to repeatedly output cuts with value $\leq (1 + O(\epsilon))\lambda$ or certify that there are no cuts of value strictly less than $(1+\epsilon)\lambda$, in $\tilde{O}(m)$ time. (The reason we output a cut with value $\leq (1 + O(\epsilon))\lambda$ is because the edge weights are maintained only approximately.) Each epoch is (conceptually) structured as follows.

- At the start of the epoch we invoke Theorem 3.1 to find $h = O(\log(n/\epsilon))$ trees $T_1, \ldots, T_h$ such that with probability $1 - \text{poly}(\epsilon/n)$, every $(1+\epsilon)$-approximate mincut of G with respect to the weights at the start of the epoch is a one-cut or two-cut of some tree T_i.
- Let $C_1, C_2, \ldots, C_r$ be an *arbitrary* enumeration of one and two-cuts induced by the trees.
- For each such cut C_j in the order, if the *current weight* of C_j is $\leq (1 + O(\epsilon))\lambda$ output it to the MWU algorithm and update weights per the MWU framework (weights only increase). Reuse C_j as long as its current weight is $\leq (1 + O(\epsilon))\lambda$.

Since weights of cuts only increase, and we consider every one and two-cut of the trees we obtain the following lemma.

Lemma 3.2. *Let $w : \mathcal{E} \to \mathbb{R}_{>0}$ be the weights at the start of the epoch and suppose mincut of G with respect to w is $\geq \lambda$; and suppose the trees $T_1, \ldots, T_h$ have the property that every $(1+\epsilon)$-approximate mincut is induced by a one or two-edge cut of one of the trees. Let $w' : \mathcal{E} \to \mathbb{R}_{>0}$ be the weights at the end of an epoch. Then, the mincut with respect to weights w' is $\geq (1+\epsilon)\lambda$.*

Within an epoch, we must consider all cuts induced by 1 or 2 edges in $\tilde{O}(1)$ spanning trees. We have complete flexibility in the order we consider them. For an efficient implementation we process the trees one at a time. For each tree T, we need to repeatedly output cuts with value $\leq (1+O(\epsilon))\lambda$ or certify that there are no cuts of value strictly less than $(1+\epsilon)\lambda$ induced by one or two edges of T. Although the underlying graph changes, we do not have to reconsider the same tree again because the cut values are non-decreasing. If each tree can be processed in $\tilde{O}(m)$ time, then an entire epoch can be processed in $\tilde{O}(m)$ time and the entire algorithm will run in $\tilde{O}\big(m/\epsilon^2\big)$ time. In Figure 3, we revise the algorithm to invoke Theorem 3.1 once per epoch and use the trees to look for approximately minimum cuts, while abstracting out the search for approximately minimum cuts in each tree.

```
held-karp-3(G = (V, E), c, ε)
  x ← 0^C,  w ← 1/c,  t ← 0,  η ← ε/ln m
  λ ← size of minimum cut w/r/t edge weights w
  while t < 1
    p ← O(log(n/ε)) spanning trees capturing every (1+ε)-apx min cut
                    w/ high probability // by Theorem 3.1
    for each tree T ∈ p
      let C_T ⊆ C be the cuts induced by removing ≤ 2 edges from T
      while (a) t < 1 and
            (b) there is a cut C ∈ C_T s.t. w̄(C) < (1+ε)λ
        y ← ⟨w, c⟩ / w̄(C) · C
        γ ← min_{e∈C} c_e
        δ ← εγ / η
        x ← x + δy
        for all e ∈ C
          w_e ← exp(εγ/c_e) · w_e
        t ← t + δ
      end while
    end for
    λ ← (1+ε)λ // start new epoch
  end while
  return x
```

Figure 3: A third implementation of the MWU framework applied to packing cuts that samples a tree packing once at the beginning of each epoch and searches for approximately minimum cuts induced by one or two edges in the sampled trees.

---

4 Cuts induced by a tree

Let $T \in \mathcal{T}$ be a rooted spanning tree of $\mathcal{G}$, and let $\lambda > 0$ be a fixed value such that the minimum cut value w/r/t w is $\geq \lambda$. We want to list cuts $C \in \mathcal{C}$ induced by 1 or 2 edges in T with weight $\overline{w}(C) \leq (1+\epsilon)\lambda$, in any order, but with a twist: each time we select one cut C, we increment the weight of every edge in C, and these increments must be reflected in all subsequent cuts. We are allowed to output cuts of value $\leq (1+O(\epsilon))\lambda$. When we finish processing T, we must be confident that there are no 1-cuts or 2-cuts $C \in \mathcal{C}_T$ induced by T with weight $\overline{w}(C) \leq (1+\epsilon)\lambda$.

The algorithmic challenge here is twofold. First, we need to find a good 1-cut or 2-cut in T quickly. Second, once a good cut is found, we have to increment the weights of all edges in the cut. Either operation, if implemented in time proportional to the number of edges in the cut, can take $O(m)$ time in the worst case. To achieve a nearly-linear running time, both operations must be implemented in polylogarithmic amortized time.

In this section we focus on finding the cuts, and assume that we can update the edge weights along any 1 or 2-cut efficiently via the `lazy-inc-cuts` data structure. Since the weights are dynamically changing, we need to describe how the algorithm finding small 1-cuts and 2-cut in T interacts with the `lazy-inc-cuts` data structure.

Interacting with the lazy weight update data structure: The interface for `lazy-inc-cuts`

Lazy Incremental Cuts: Interface

Operations	Description
`lazy-inc-cuts`	Given an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with capacities $c : \mathcal{E} \to \mathbb{R}_{>0}$ and edge weights $w : \mathcal{E} \to \mathbb{R}_{>0}$, and a rooted spanning tree T, initializes the `lazy-inc-cuts` data structure.
`inc-cut`	Given a 1-cut or 2-cut $C \in \mathcal{C}_T$ induced by T, simulates a weight increment along the cut per the MWU framework. Returns a list of tuples $\Delta = \{(e_1, \delta_1), (e_2, \delta_2), \ldots\}$, where each $e_i \in \mathcal{E}$ is an edge and each $\delta_i > 0$ is a positive real value that signals that the weight of the edge e_i has increased by an additive factor of δ_i. Each increment (e_i, δ_i) satisfies $\delta_i \geq \Omega\big(\epsilon w_e / \log^2 n\big)$.
`flush()`	Returns a list of tuples $\Delta = \{(e_1, \delta_1), \ldots, (e_k, \delta_k)\}$ of all residual weight increments not accounted for in previous calls to `inc-cut` and `flush`.

Figure 4: The interface for the data structure `lazy-inc-cuts`, which efficiently simulates weight increments along cuts induced by 1 or 2 edges in a fixed and rooted spanning tree. Implementation details are provided later in Section 5.

---

is given in Figure 4. The primary function of the `lazy-inc-cuts` data structure is `inc-cut`, which takes as input one or two edges in T, and simulates a weight update along the corresponding 1-cut or 2-cut. `inc-cut` returns a list of weight increments (e_i, δ_i), where each e_i is an edge and $\delta_i > 0$ is a positive increment that we actually apply to the underlying graph. The cumulative edge weight increments returned by the subroutine will always underestimate the true increase, but by at most a $(1 + O(\epsilon))$-multiplicative factor for every edge. That is, if we let w_e denote the "true weight" of an edge e implied by a sequence of weight increments along cuts, and $\widetilde{w}_e$ the sum of δ increments for e returned by `inc-cut`, then the `lazy-inc-cuts` data structure maintains the invariant that $\widetilde{w}_e \leq w_e \leq (1 + \rho\epsilon)\widetilde{w}_e$, where ρ is a fixed constant that we can adjust. In particular, a cut of weight $\leq (1+\epsilon)\lambda$ w/r/t the approximated edge weights retains weight $\leq (1+O(\epsilon))\lambda$ in the true underlying graph. Each increment (e, δ) returned by `inc-cut` reflects a $(1+\Omega(\epsilon/\log^2 n))$-multiplicative increase to w_e. As the weight of an edge is bounded above by $\tilde{O}\big(m^{1/\epsilon}\big)$, the subroutine returns an increment for a fixed edge e at most $O\big(\log^3 n/\epsilon^2\big)$ times total over the entire algorithm.

A second routine, `flush()`, returns a list of increments that captures all residual weight increments not accounted for in previous calls to `inc-cut` and `flush`. The increments returned by `flush()` may not be very large, but restores $\widetilde{w}_e = w_e$ for all e. Formally, we will refer to the following.

Lemma 4.1. *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be an undirected graph with positive edge capacities $c : \mathcal{E} \to \mathbb{R}_{>0}$ and positive edge weights $w : \mathcal{E} \to \mathbb{R}_{>0}$, and T a rooted spanning tree in $\mathcal{G}$. Consider an instance of* `lazy-inc-cuts(`$\mathcal{G}$`,`c`,`w`,`T`)` *over a sequence of calls to* `inc-cut` *and* `flush`*. For each edge e, let w_e be the "true" weight of edge e if the MWU framework were followed exactly, and let $\widetilde{w}_e$ be the approximate weight implied by the weight increments returned by calls to* `inc-cut` *and* `flush`*.*

(a) After each call to `inc-cut`*, we have $\widetilde{w}_e \leq w_e \leq (1 + O(\epsilon))\widetilde{w}_e$ for all edges $e \in \mathcal{E}$.*

(b) After each call to `flush`*, we have $\widetilde{w}_e = w_e$.*

(c) Each call to `inc-cut` *takes $O(\log^2 n + I)$, where I is the total number of weight increments returned by* `inc-cut`*.*

(d) Each call to `flush` *takes $O(m \log^2 n)$ time and returns at most m weight increments.*

```
held-karp-4(G = (V, E), c, ε)
  x ← 0^C,  w ← 1^E,  t ← 0,  η ← ε/ln m
  λ ← size of minimum cut w/r/t edge weights w
  while t < 1
    p ← O(log(n/ε)) spanning trees capturing all (1+ε)-apx min cuts
      w/r/t weights w w/ high probability // by Theorem 3.1
    for each tree T ∈ p
      lic ← lazy-inc-cuts.init(G, c, T)
      let C_T ⊆ C be the cuts induced by removing ≤ 2 edges from T
      while (a) t < 1 and until
            (b) we have certified there are no cuts C ∈ C_T
                s.t. w̄(C) < (1+ε)λ
        find C ∈ C_T w/ w̄(e) < (1+O(ε))λ
        y ← (⟨w, c⟩ / w̄(C)) · C
        γ ← min_{e∈C} c_e
        δ ← εγ/η
        x ← x + δy
        t ← t + δ
        Δ ← lic.inc-cut(C)
        for (e, δ) ∈ Δ
          w_e ← w_e + δ
        end for
      end while
      Δ ← lic.flush()
      for (e, δ) ∈ Δ
        w_e ← w_e + δ
      end for
    end for
  end while
  return x
```

Figure 5: A fourth implementation of the MWU framework applied to the Held-Karp relaxation, that integrates the `lazy-inc-cuts` data structure into the subroutine that processes each tree.

(e) Each increment (e, δ) *returned by* `inc-cut` *satisfies* $\delta \geq \Omega\big((\epsilon/\log^2 n)w_e\big)$.

The implementation details and analysis of `lazy-inc-cuts` are given afterwards in Section 5. An enhanced sketch of the algorithm so far, integrating the `lazy-inc-cuts` data structure but abstracting out the search for small 1- and 2-cuts, is given in Figure 5. In this section, we assume Lemma 4.1 and prove the following.

Lemma 4.2. *Let* T *be a rooted spanning tree and* $\lambda > 0$ *a target cut value. One can repeatedly output 1-cuts and 2-cuts* $C \in \mathcal{C}_T$ *induced by* T *of weight* $\overline{w}(C) \leq (1 + O(\epsilon))\lambda$ *and increment the corresponding edge weights (per the MWU framework) until certifying that there are no 1-cuts or 2-cuts* $C \in \mathcal{C}_T$ *of value* $\overline{w}(C) \leq (1+\epsilon)\lambda$ *in* $O\big(m \log^2 n + K \log^2 n + I \log n\big)$ *time, where* K *is the number of 1-cuts and 2-cuts of value* $\leq (1 + O(\epsilon))\lambda$ *output and* I *is the total number of weight*

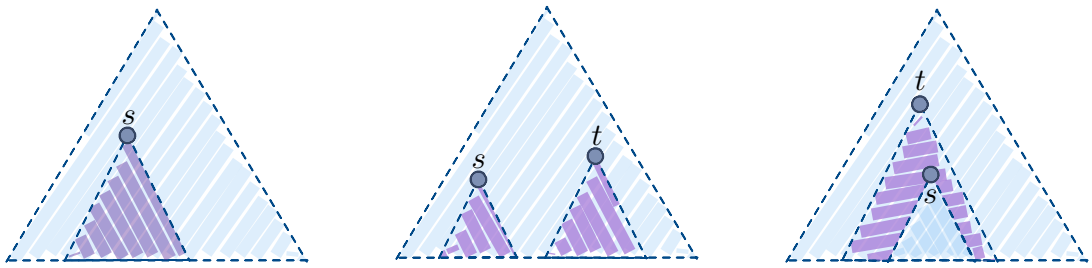

Figure 6: 3 types of tree-cuts

increments returned by the `lazy-inc-cuts` *data structure.*

The rest of the section is devoted to the proof of the preceding lemma. We first set up some notation. Fixing a rooted spanning tree T induces a partial order on the vertices $\mathcal{V}$. For two vertices u and v, we write $u \leq v$ if u is a descendant of v. Two vertices u and v are *incomparable*, written $u \parallel v$, if neither descends from the other. For each vertex $v \in \mathcal{V}$, let T_v be the subtree of T rooted at v. Let $D(v) \stackrel{\text{def}}{=} \mathcal{V}(T_v) = \{u : u \leq v\}$ be the *down set* of all descendants of v in T (including v). We work on a graph with edge weights given by a vector $w : \mathcal{E} \to \mathbb{R}_{>0}$. For a (proper) set of vertices $S \subset \mathcal{V}$, we let $\mathcal{C}(S) \stackrel{\text{def}}{=} \{e \in \mathcal{E} : |e \cap S| = 1\}$ be the set of edges cut by S, and let $\overline{\mathcal{C}}(S) \stackrel{\text{def}}{=} \overline{w}(\mathcal{C}(S)) = \sum_{e \in \mathcal{C}(S)} w_e$ denote the weight of the cut induced by S. For two disjoint sets of vertices $S, T \subseteq \mathcal{V}$, we let $\mathcal{C}(S,T) = \{e = (a,b) \in \mathcal{E} : a \in S, b \in T\}$ be the edges crossing from S to T, and let $\overline{\mathcal{C}}(S,T) = \overline{w}(\mathcal{C}(S,T))$ be their sum weight.

Per Karger [2000], we divide the 1-cuts and 2-cuts into three distinct types, drawn in Figure 6. First there are the 1-cuts of T of the form $\mathcal{C}(D(s))$. Then we have 2-cuts $\mathcal{C}(D(s) \cup D(t))$ where s and t are incomparable (i.e., $s \parallel t$). We call these *incomparable 2-cuts*. Lastly, we have 2-cuts $\mathcal{C}(D(t) \setminus D(s))$ where s is a descendant of t (i.e., $s < t$). We call these *nested 2-cuts*. For a fixed tree, we process all 1-cuts, incomparable 2-cuts, and then nested 2-cuts, in this order. For each of these three cases we prove a lemma similar to Lemma 4.2 but restricted to the particular type of cut. Throwing in $\tilde{O}(m)$ time to initialize the `lazy-inc-cuts` data structure at the beginning and $\tilde{O}(m)$ to call `flush` at the end then gives Lemma 4.2.

Standard data structures. Following Karger [2000], we rely on the dynamic tree data structure, *link-cut trees*, of Sleator and Tarjan [1983]. Link-cut trees store real values at the nodes of a rooted tree T and support aggregate operations along node-to-root paths in logarithmic time. We employ two operations in particular. First, given a node v and a value $\alpha \in \mathbb{R}$, we can add α to the value of every node on the v-to-root path (i.e., every node u such that $u \geq v$) in $O(\log n)$ amortized time. Second, given a node v, we can find the minimum value of any node on the v-to-root path in $O(\log n)$ amortized time.

We also need to compute least common ancestors (`lca`) in T for pairs of nodes. Harel and Tarjan [1984] showed how to preprocess T in linear time and answer `lca` queries in constant time.

4.1 1-cuts

Let $e = (t, s) \in T$ be a tree edge with t a parent of s. Then $T - e$ consists of two components: the subtree T_s rooted at s, and the remaining tree $T \setminus T_s$. The cut $\mathcal{C}(D(s))$ in the underlying graph $\mathcal{G}$ induced by the subtree T_s is a "1-cut" of T. In this section, we want to output a sequence of 1-cuts

$C \in \mathcal{C}_T$ in T with $\overline{w}(C) \leq (1 + O(\epsilon))\lambda$, repeating if necessary, while incrementing weights of each listed cut on the fly, until we've determined that there are no 1-cuts C with weight $\overline{w}(C) \leq (1+\epsilon)\lambda$.

Computing and maintaining the 1-cut values. To find the minimum 1-cut induced by a down set $D(v)$, Karger observes that the weight $\overline{\mathcal{C}}(D(v))$ can be rewritten as

$$\overline{\mathcal{C}}(D(v)) = \sum_{x \in D(v)} \overline{\mathcal{C}}(x) - 2 \sum_{e \in \mathcal{E}[D(v)]} w_e,$$

where $\overline{\mathcal{C}}(x)$ is the weighted degree of the vertex x, and $\mathcal{E}[D(v)]$ is the set of edges with both endpoints in $D(v)$.

The quantities $\sum_{x \in D(v)} \overline{\mathcal{C}}(x)$ over $v \in \mathcal{V}$ are just tree sums over the weighted degrees, and Karger [2000] computes them once for all v in $O(m)$ total time by a depth-first traversal. In the dynamic setting, we can compute and maintain the tree sums in $O(\log n)$ time per edge update in a link-cut tree. Each time we increment the value of an edge $e = (u, v) \in \mathcal{E}$ by some $\delta > 0$, we add δ to the value of every node on the paths to the root from u and v.

The sum $\overline{w}(\mathcal{E}[D(v)]) \stackrel{\text{def}}{=} \sum_{e \in \mathcal{E}[D(v)]} w_e$ over edges with both endpoints in $D(v)$ is just the sum of weights of all edges $e = (x, y)$ for which the least common ancestor $\texttt{lca}(x, y)$ of x and y is in $D(v)$; i.e., $\overline{w}(\mathcal{E}[D(v)]) = \sum_{e:\texttt{lca}(e) \in D(v)} w_e$ By Harel and Tarjan [1984], we can compute $\texttt{lca}(e)$ for every edge $e \in \mathcal{E}$ in linear time. We then compute for each vertex v the sum $\sum_{\texttt{lca}(e)=v} w_e$ of all weights of edges whose least common ancestor is v. The quantities $\overline{w}(\mathcal{E}[(D(v))])$ can then be computed as tree-sums by depth-first search, and updated dynamically with link-cut trees, similar to before.

The preceding discussion shows that the sums $\overline{\mathcal{C}}(D(v))$ for each v can be maintained dynamically in logarithmic time using link-cut trees. This is summarized in the lemma below.

Lemma 4.3. *Given edge-weighted graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ and a rooted spanning tree T, there is a data structure to maintain $\overline{\mathcal{C}}(D(v))$ for all $v \in \mathcal{V}$ as edge-weights are changed. The initialization cost is $O(m)$, updates and queries take $O(\log n)$ amortized time.*

Processing all 1-cuts. The weight of all 1-cuts can be computed in $O(m)$ time statically and $O(\log n)$-time per edge update. To process all 1-cuts of the tree, we consider all downsets $D(v)$ in any order. For each vertex v, we check if $\overline{\mathcal{C}}(D(v)) \leq (1 + \epsilon)\lambda$ in $O(\log n)$ amortized time. If not, then we move on to the next vertex. Since the edge weights are monotonically increasing, $\overline{\mathcal{C}}(D(v))$ will never be $\leq (1+\epsilon)\lambda$ again, so we do not revisit v. Otherwise, if $\mathcal{C}(D(v))$ is good, we take the cut in our solution and pass v to `inc-cut` to signal a weight increment along $\mathcal{C}(D(v))$. In turn, `inc-cut` returns a list of "official" edge increments $\{(e_i, \delta_i)\}$. Each update can be incorporated into the link-cut tree in logarithmic time by Lemma 4.3. After processing the edge increments returned by `inc-cut`, we continue to process $D(v)$ until $\mathcal{C}(D(v))$ is no longer a good cut. The total running time to process all 1-cuts is $O(m)$ plus $O\big(\log^2 n\big)$ per cut output and $O(\log n)$ per edge weight increment returned by `inc-cut`.

Lemma 4.4. *Let T be a fixed and rooted spanning tree and $\lambda > 0$ a fixed target cut value. Employing the* `inc-cut` *routine of the* `lazy-inc-cuts` *data structure to approximately increment edge weights, one can repeatedly find 1-cuts C induced by T of value $\overline{\mathcal{C}}(C) \leq (1 + O(\epsilon))\lambda$ and increment the corresponding edge weights (per the MWU framework) until certifying that there are no 1-cuts C of value $\overline{\mathcal{C}}(C) \leq (1 + \epsilon)\lambda$ in $O\big(m + K \log^2 n + I \log n\big)$ total running time, where K is the number of 1-cuts of value $\leq (1+O(\epsilon))$ output and I is the total number of weight increments returned by* `inc-cut`.

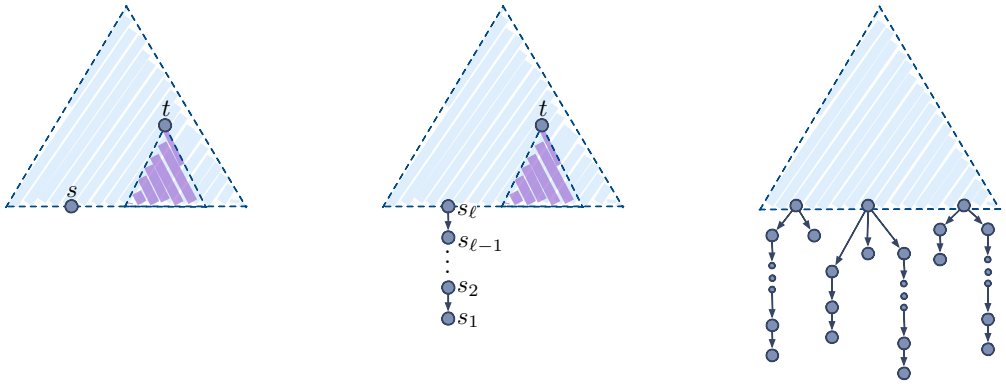

Figure 7: 3 cases for incomparable 2-cuts

4.2 2-cuts on incomparable vertices

In this section, we consider incomparable 2-cuts of the form $\mathcal{C}(D(s) \cup D(t))$, where $s \parallel t$. We first consider the case where s is a fixed leaf, and t ranges over all vertices incomparable to s. We then extend the leaf case to consider the case where s ranges over a path in the tree down to a leaf, and t ranges over all incomparable vertices to the path. We then apply an induction step that reduces processing the whole tree to a logarithmic number of rounds where in each round we process all the paths to leaves. See Figure 7.

Incomparable 2-cuts with a leaf. We first fix s to be a leaf in T, and consider all incomparable 2-cuts with one side fixed to be $\{s\}$. Karger observes that the weight of the cut induced by $D(t)+s$, for $t \parallel s$, can be written as

$$\overline{\mathcal{C}}(D(t) + s) = \overline{\mathcal{C}}(s) + \overline{\mathcal{C}}(D(t)) - 2\overline{\mathcal{C}}(D(t), s)$$

where $\overline{\mathcal{C}}(D(t), s)$ is the weight of edges with one endpoint s and the other in $D(t)$. We only need to consider down sets $D(t)$ for which s and $D(t)$ are connected by some edge, since otherwise either $\{s\}$ or $D(t)$ induce a smaller cut, and after processing all 1-cuts, neither $\{s\}$ or $D(t)$ induce small enough cuts.

For fixed s and ranging t over all $t \parallel s$, Karger finds the weight of the best incomparable $D(t) + s$ cut statically as follows. As $\overline{\mathcal{C}}(s)$ is fixed, the t's are differentiated only by the values $\overline{\mathcal{C}}(D(t)) - 2\overline{\mathcal{C}}(D(t), s)$. Karger's algorithm creates a dynamic tree over T with each vertex t initialized with the value $\overline{\mathcal{C}}(D(t))$. (This initial tree only needs to be constructed once, and immutably reused for other s's). Karger's algorithm immediately adds $+\infty$ to the value of every vertex on the s-to-root path to eliminate all comparable vertices from consideration. For each edge $e = (s, v)$ with endpoint s, Karger's algorithm subtracts $2w_e$ from every vertex on the v-to-root path, to incorporate the second term $-2\overline{\mathcal{C}}(D(t), s)$. Then, for each edge $e = (s, v)$ adjacent to s, Karger's algorithm finds the minimum value vertex on the v to root path. The minimum value vertex z over all such queries gives the best $s + D(z)$ cut subject to $s \parallel z$, as desired. The running time to process the leaf s, excluding the construction of the initial tree, is bounded by the time spent subtracting edge weights and taking minimums over the v-to-root path for each edge (s, v) incident to s. With dynamic trees, subtracting and finding the minimum along a vertex-to-root path each take $O(\log n)$ amortized time. Thus, up to a logarithmic factor, the time spent processing all incomparable 2-cuts with one side a fixed leaf s is proportional to the number of edges incident to s.

In the dynamic setting, we do the following. We first compute the sum $\overline{\mathcal{C}}(D(t)) - 2\overline{\mathcal{C}}(D(t), s)$ for all incomparable $t : t \parallel s$, as above. We then run through the edges (s, v) incident to s in any fixed order, checking for $s + D(t)$ cuts with $t \geq v$ that have value $\leq (1+\epsilon)\lambda$ and updating edge weights along good cuts on the fly, as follows.

Fix an edge $e = (s, v)$ incident to s. We process e until we know that there is no good $(D(t)+s)$-cut for any ancestor t of v. We first find the minimum value vertex on the v-to-root path in $O(\log n)$ amortized time. If the minimum value is $> (1+\epsilon)\lambda - \overline{\mathcal{C}}(D(s))$, then we move onto the next edge incident to s. Otherwise, let t be an ancestor of v such that $s+D(t)$ induces a good cut. We pass the cut $\mathcal{C}(s + D(t))$ (in terms of the roots of s and t) to `inc-cut`, which simulates a weight update along $\mathcal{C}(s + D(t))$ and returns a list of weight increments $\{(e_i, \delta_i)\}$. For each such edge $e = (x, y)$ whose weight is increased by some $\delta > 0$, we add δ to every vertex on the x-to-root and y-to-root paths. We then subtract 2δ from the $\texttt{lca}(x, y)$-to-root path. These additions and subtraction account for the first term, $\overline{\mathcal{C}}(D(t))$. If $x = s$, then we also subtract 2δ from the y-to-root path for the sake of the second term, $-2\overline{\mathcal{C}}(D(t), s)$. After all official weight increments returned by `inc-cut` have been incoporated into the dynamic tree, we continue to process the same edge $e = (s, v)$. Once all edges $e = (s, v)$ incident to s have been processed, we have certified that all 2-cuts with one side fixed to the leaf s have weight $\geq (1+\epsilon)\lambda$. The total time is $O(\log n)$ times the number of edges incident to s, plus $O\big(\log^2 n\big)$ time for every cut output and $O(\log n)$ for every increment returned by `inc-cut`.

Leaf-paths. The approach outlined above for processing leaves extends to processing paths. Let a *leaf-path* be a maximal subtree that is a path. Let $p = s_1, s_2, \ldots, s_\ell$ be a leaf path, where s_1 is a leaf, s_2 is its parent, and so forth. We want to process all 1-cuts of the form $D(s_i) \cup D(t)$, where t is incomparable to any and all of the s_i's on the path.

For each i, we have

$$\overline{\mathcal{C}}(D(s_i) \cup D(t)) = \overline{\mathcal{C}}(D(s_i)) + \overline{\mathcal{C}}(D(t)) - 2\overline{\mathcal{C}}(D(s_i), D(t)),$$

where $\overline{\mathcal{C}}(D(s_i), D(t))$ is the weight of all edges between $D(s_i)$ and $D(t)$. The values $\overline{\mathcal{C}}(D(t))$ for $t \parallel s_1, \ldots, s_\ell$ and $\overline{\mathcal{C}}(D(s_i))$ for each i are easy to compute initially and maintain dynamically, per the earlier discussion about 1-cuts in Section 4.1. For fixed i, the incomparable t are differentiated by the value $\overline{\mathcal{C}}(D(t)) - 2\overline{\mathcal{C}}(D(s_i), D(t))$. The basic idea here is that $\overline{\mathcal{C}}(D(s_{i+1}), D(t))$ can be written in terms of the preceding cut in the form $\overline{\mathcal{C}}(D(s_{i+1}), D(t)) = \overline{\mathcal{C}}(D(s_i), D(t)) + \overline{\mathcal{C}}(s_{i+1}, D(t))$. That is, after computing $\overline{\mathcal{C}}(D(s_i), D(t))$ for all incomparable t, we can keep these values and just incorporate the weights of edges incident to s_{i+1} to get the values $\overline{\mathcal{C}}(D(s_i), D(t))$ for all incomparable t.

We first review the static case, where we want to find the minimum such cut. Karger's algorithm first processes the leaf s_1 as a leaf, as outlined above, and takes note of the minimum $s_1 + D(t)$ cut. It then processes s_2 in the same fashion, except it continues the aggregate data built when processing s_1 rather than starting fresh. By subtracting the weights of all edges incident s_2 from the values of the various incomparable downsets, and having already subtracted the weights of edges incident to s_1, we will have subtracted the weight of all edges incident to $D(s_2) = \{s_1, s_2\}$, and so we have computed $\overline{\mathcal{C}}(D(t)) - 2\overline{\mathcal{C}}(D(s_2), D(t))$ for all $t \parallel s_2$. In this fashion, Karger's algorithm marches up the path doing essentially the same work for each vertex as for a fixed leaf, except continuing the same link-cut tree from one s_i to the next. The total work processing the entire leaf-path is, up to a logarithmic factor, proportional to the total number of edges incident to any node s_i on the path.

The analogous adjustments are made in the dynamic setting. That is, we process s_1 as a leaf, passing good cuts to `inc-cut` and incorporating the returned edge increments on the fly. After processing s_1, we keep the accumulated aggregate values, and start processing s_2 likewise. Marching up the leaf-path one vertex at a time. In this manner, we are able to process 2-cuts of the form

$D(s_i) + D(t)$ over each s_i along the leaf-path $s_1, \ldots, s_\ell$ (subject to $t \parallel s_1, \ldots, s_\ell$) and at the end certify there are no such cuts of weight $< (1+\epsilon)\lambda$. The total work in processing the leaf path is $O(\log n)$ for each edge incident to $s_1, \ldots, s_\ell$, $O\big(\log^2 n\big)$ for each good cut output, and $O(\log n)$ for each weight increment returned from the subroutine `inc-cut`.

All incomparable 2-cuts. Karger [2000] showed that an efficient subroutine for processing all incomparable 2-cuts on a single leaf-path leads to an efficient algorithm for processing all incomparable 2-cuts in the entire tree as follows. The overall algorithm processes the 2-cuts of a tree in phases. Each phase processes all the leaf-paths of the current tree, and then contracts the leaf-paths into their parents and recurses on the new tree in a new phase. Each leaf in the contracted graph had at least two children in the previous phase, so the number of nodes has halved. After at most $O(\log n)$ phases, we have processed all the incomparable 2-cuts in the tree. Excluding the work for selecting a cut and incrementing edge weights returned by the `lazy-inc-cuts` data structure, processing a leaf-path takes time proportional to the number of edges incident to the leaf-path. Each edge is incident to at most 2 leaf-paths. Thus, a single phase takes $O(m \log n)$ time, in addition to $O\big(\log^2 n\big)$ time for every cut output and $O(\log n)$ amortized time for every edge weight increment returned by the `lazy-inc-cuts` data structure. In conclusion, we have shown the following.

Lemma 4.5. *Let T be a fixed and rooted spanning tree and $\lambda > 0$ a fixed target cut value. Employing the* `inc-cut` *routine of the* `lazy-inc-cuts` *data structure to approximately increment edge weights, one can repeatedly find incomparable 2-cuts C induced by T of value $\overline{\mathcal{C}}(C) \leq (1+O(\epsilon))\lambda$ and increment the corresponding edge weights (per the MWU framework) until certifying that there are no incomparable 2-cuts C of value $\overline{\mathcal{C}}(C) \leq (1+\epsilon)\lambda$ in $O\big(m \log^2 n + K \log^2 n + I \log n\big)$ total amortized time, where K is the number of incomparable 2-cuts of value $\leq (1+O(\epsilon))\lambda$ output and I is the total number of weight increments returned by* `inc-cut`.

4.3 Nested 2-cuts

The other type of 2-cuts is of the form $D(t) \setminus D(s)$, where t is an ancestor of s. We call these cuts *nested 2-cuts*. As in the case of incomparable 2-cuts, we first consider the case where s is fixed to be a leaf, then the case where s is along a leaf-path, and then finally an induction step that processes all leaf-paths in each pass.

Nested 2-cuts with a leaf. We take the same approach as with incomparable 2-cuts, and start with the case where s is a fixed leaf. Karger observed that the weight of the cut induced by $D(t) - s$ can be written as

$$\overline{\mathcal{C}}(D(t) - s) = \overline{\mathcal{C}}(D(t)) + \overline{\mathcal{C}}(s, D(t)),$$

where $\overline{\mathcal{C}}(s, D(t))$ is the sum weight of the edges between s and $D(t)$. We start with a link-cut tree over T where each node v is initialized with $\overline{\mathcal{C}}(D(v))$. For each edge $e = (s, u)$ incident to s, we add w_e to every node on the u-to-root path in $O(\log n)$ time. In the static setting, Karger's algorithm finds the minimum value on the s-to-root path to get the value of the minimum $(D(t) - s)$-cut over all $t > s$. In the dynamic setting, we repeatedly find the minimum value on the s-to-root path as long as this value is below the threshold $(1+\epsilon)\lambda$. Each time we find a good cut $\mathcal{C}(D(t) - s)$, we send the cut to the routine `inc-cut` (compactly represented by s and t). For each edge weight increment $(e = (x, y), \delta)$ returned by `inc-cut`, we add δ to every the value of every vertex on the x-to-root and y-to-root paths and subtract 2δ from every value of every vertex on the $\texttt{lca}(x, y)$ path in $O(\log n)$ time, to account for the first term $\overline{\mathcal{C}}(D(t))$ (see Section 4.1). If e is also incident to s, then we add δ to the value of every vertex on the s-to-root path to correct the second term

$\overline{\mathcal{C}}(s, D(t))$. The total amortized running time for processing all nested 2-cuts $D(t) \setminus s$ with a fixed leaf s is $O(\log n)$ times the number of edges incident to s, plus $O\big(\log^2 n\big)$ for every cut output and $O(\log n)$ for every edge weight increment returned by the `inc-cut` subroutine.

Nested 2-cuts along a leaf-path. We extend the leaf case to a leaf path $s_1, \ldots, s_\ell$, where s_1 is a leaf and each s_i with $i > 1$ has exactly one child, and consider all cuts of the form $D(t) \setminus D(s_i)$ over all s_i in the leaf path and all ancestors t of s_i. For each i, we can write

$$\begin{aligned}\overline{\mathcal{C}}(D(t) \setminus D(s_i)) &= \overline{\mathcal{C}}(D(t)) + \overline{\mathcal{C}}(D(s_i), D(t)) \\ &= \overline{\mathcal{C}}(D(t)) + \sum_{j \le i} \overline{\mathcal{C}}(s_j, D(t)) - 2 \sum_{e \in \mathcal{E}[D(s_i)]} w_e,\end{aligned}$$

where $\mathcal{E}(D(S_i))$ is the set of edges with both endpoints in $D(s_i)$. When $i = 1$, this is the same equation we obtained for the leaf case. The first term, $\overline{\mathcal{C}}(D(t))$, is just the weight of the 1-cut $D(t)$, and easy to compute statically and maintain dynamically per the earlier discussion in Section 4.1. The second term, $\sum_{j \le i} \overline{\mathcal{C}}(s_j, D(t))$, can be computed from that of s_{i-1} by incorporating the edge weights of every edge incident to s_i. The last term, $\sum_{e \in \mathcal{E}[D(s_i)]} w_e$, is independent of t and easy to compute and maintain dynamically for all i, per the discussion for a similar term that arose in computing 1-cuts in Section 4.1.

These observations lead to the following bottom-up approach. We first process s_1 as though a leaf, by incrementing values on x-to-root paths for each edge (s_1, x) incident to s_1, and then looking for small nested 2-cuts $D(t) - s_1$ by checking the minimum value on the s_1-to-root path, incrementing weights along good $D(t) - s_1$ min-cuts as they are found via the `lazy-inc-cuts` data structure. Once we have certified that are no nested $D(t) - s_1$ cuts of value $< (1 + \epsilon)\lambda$, we process s_2 likewise, except continuing the aggregate values computed when processing s_1. For general $i > 1$, we find good nested 2-cuts of the form $D(t) \setminus D(s_i)$ by retaining the aggregate values computed previously from looking for good $D(t) \setminus D(s_{i-1})$ cuts, adding the weights of edges incident to s to fix the $\overline{\mathcal{C}}(D(s_i), D(t))$ term, and then repeatedly updating edge weights along good nested $(D(t) \setminus D(s_i))$-cuts via `inc-cut`. As with incomparable 2-cuts along leaf-paths, the total amortized running time is $O(\log n)$ times the number of edges incident to $s_1, \ldots, s_\ell$, plus $O\big(\log^2 n\big)$ work for each cut output and $O(\log n)$ work for each edge increment returned by `inc-cut`.

All nested 2-cuts. By the same induction step as for incomparable 2-cuts, efficiently processing leaf-paths leads to a procedure for efficiently processing the whole tree. The processing is broken into phases. Each phase processes all the leaf-paths of the current tree, and then contracts the leaf-paths into their parents. Each leaf in the contracted tree had at least two children previously, so the number of nodes has at least halved. After $O(\log n)$ phases, we have processed the entire tree. Modulo $O\big(\log^2 n\big)$ work for each cut output and $O(\log n)$ work for each edge increment returned by `inc-cut`, processing a leaf path takes $O(\log n)$ time for each edge incident to any node on the path, and conversely each edge is incident to at most 2 leaf paths. Thus, excluding the time spent outputting cuts and incrementing edge weights along the cuts, each phase takes $O(m \log n)$ to complete.

Lemma 4.6. *Let T be a fixed and rooted spanning tree and $\lambda > 0$ a fixed target cut value. Employing the* `inc-cut` *routing of the* `lazy-inc-cuts` *data structure to manage edge weight increments, one can repeatedly find nested 2-cuts C induced by T of value $\overline{\mathcal{C}}(C) \le (1 + O(\epsilon))\lambda$ and increment the corresponding edge weights (per the MWU framework) until certifying that there are no incomparable 2-cuts C of value $\overline{\mathcal{C}}(C) \le (1 + \epsilon)\lambda$, in $O\big(m \log^2 n + K \log^2 n + I \log n\big)$ total running time, where K is the number of nested 2-cuts of value $\le (1 + O(\epsilon))\lambda$ output and I is the total number of weight increments returned by* `inc-cut`.

5 Applying multiplicative weight updates along cuts

We address the remaining issue of implementing the `lazy-inc-cuts` data structure for a fixed and rooted spanning tree T. An interface for `lazy-inc-cuts` was given previously in Figure 4 and we target the bounds claimed by Lemma 4.1. Recall that the MWU framework takes a cut $C \in \mathcal{C}$, computes the smallest capacity $\gamma = \arg\min_{e \in C} c_e$ in the cut, and increases the weight of each cut-edge $e \in C$ by a multiplicative factor of $\exp(\epsilon\gamma/c_e)$ (see Figure 1). A subtle point is that the techniques of Section 4 identify approximate min-cuts without explicitly listing the edges in the cut. A 1-cut $\mathcal{C}(D(s))$ is simply identified by the root s of the down-set $D(s)$, and likewise 2-cuts of the form $\mathcal{C}(D(s) \cup D(t))$ (when $s \parallel t$) and $\mathcal{C}(D(t) \setminus D(s))$ (when $s < t$) can be described by the two nodes s and t. In particular, an approximately minimum cut C is identified without paying for the number of edges $O(|C|)$ in C in the running time. When it comes to updating the edge weights, the natural approach of visiting each edge $e \in C$ would be too slow, to say nothing of even identifying all the edges in C in time proportional to $|C|$.

While the general task of incrementing[2] weights along a cut appears difficult to execute both quickly and exactly, we have already massaged the setting to be substantially easier. First, we can afford to approximate the edge weights w_e by a small multiplicative error. This means, for example, that a cut edge $e \in C$ with very large capacity $c_e \gg \gamma$ can to some extent be ignored. Second, we are not incrementing weights along any cut, but just the 1-cuts and 2-cuts of a fixed rooted spanning tree T. We have already seen in Section 4 that restricting ourselves to 1-cuts and 2-cuts allows us to (basically) apply dynamic programming to find small cuts, and also allows us to use dynamic trees to efficiently update and scan various values in the aggregate. Here too we will see that 1-cuts and 2-cuts are simple enough to be represented efficiently in standard data structures. More precisely, we will show that the 1-cuts and 2-cuts $\mathcal{C}_T$ have the following nice structure.

Lemma 5.1. *Let T be a fixed rooted spanning tree of an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with $|\mathcal{E}| = m$ edges and $|\mathcal{V}| = n$ vertices. In $O\big(m \log^2 n\big)$ time, one can construct a collection of nonempty cuts $\mathcal{D}_T \subseteq 2^{\mathcal{E}}$ such that*

(i) every edge $e \in \mathcal{E}$ appears in at most $O(\log^2 n)$ cuts $D \in \mathcal{D}_T$, and

(ii) every 1-cut or 2-cut $C \in \mathcal{C}_T$ (described succinctly by at most 2 roots of subtrees) can be decomposed into the disjoint union $C = D_1 \sqcup \cdots \sqcup D_\ell$ of $\ell = O(\log^2 n)$ cuts $D_1, \ldots, D_\ell \in \mathcal{D}_T$ in $O(\log^2 n)$ time.

By building the collection $D \in \mathcal{D}_T$ once for a tree T, Lemma 5.1 reduces the problem of incrementing along any 1-cut or 2-cut $C \in \mathcal{C}_T$ to incrementing along a "canonical" cut $D \in \mathcal{D}_T$ known *a priori*. This is important because multiplicative weight updates can be applied to a *static* set relatively efficiently by known techniques [Young, 2014, Chekuri and Quanrud, 2017]. It is also important that cuts in $\mathcal{D}_T$ are sparse, in the sense that $\sum_{D \in \mathcal{D}_T} |D| = O\big(m \log^2 n\big)$, as this sum factors directly into the running time guarantees. We first prove Lemma 5.1 in Section 5.1, and then in Section 5.2 we show how to combine amortized data structures for each $D \in \mathcal{D}_T$ to increment along any 1-cut or 2-cut $C \in \mathcal{C}_T$.

5.1 Canonical cuts

[2] In the MWU framework the weights increase in a multiplicative fashion. We use the term "incrementing" in place of "updating" since weights are increasing and also because it is convenient to think in terms of the logarithm of the weights, which do increase additively.

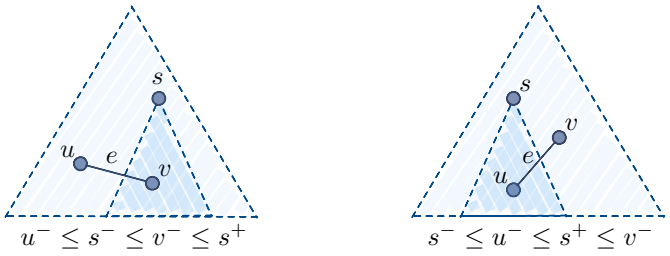

Figure 8: An edge cut e by a subtree T_s, and two possible inequalities in the Euler order.

Euler tours and orderings. Let T be a fixed and rooted tree on n vertices $\mathcal{V}$, and fix an Euler tour on a bidirected copy of T that replaces each edge of T with arcs in both directions, starting from the root. For each vertex v, we create two symbols: v^- means we enter the subtree T_v rooted at v, and v^+ means we leave the subtree v. Let $\mathcal{V}^\pm = \{v^-, v^+ : v \in \mathcal{V}(T)\}$ denote the whole collection of these $2n$ symbols. The Euler tour enters and leaves each subtree exactly once in a fixed order. Tracing the Euler tree induces a unique total ordering on $\mathcal{V}^\pm$ (see the picture on the right).

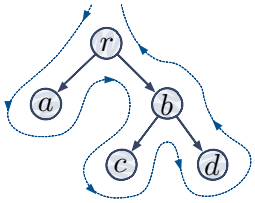

An Euler tour inducing the ordering $r^- < a^- < a^+ < b^- < c^- < c^+ < d^- < d^+ < b^+ < r^+$.

This ordering has a couple of interesting properties. For every vertex $v \in \mathcal{V}(T)$, we have $v^- < v^+$. Letting r denote the root of T, r^- is the first element in the ordering and r^+ is the last element in the ordering. More generally, for a vertex v and a vertex $w \in D(v)$ in the subtree rooted at v, we have $v^- \le w^- < w^+ \le v^+$, with all inequalities strict if $v \neq w$. That is, the ranges $\{[v^-, v^+] : v \in \mathcal{V}(T)\}$ between entering and leaving a subtree form a laminar set.

The Euler order on the vertices endows a sort of geometry to the edges. Each edge $e = (u, v)$ (with $u^- < v^-$) can be thought of as an interval $[u^-, v^-]$. A downset $D(s)$ cuts e iff $u^- \le s^- \le v^- \le s^+$ or $s^- \le u^- \le s^+ \le v^-$ (see Figure 8); that is, iff $|\{u^-, v^-\} \cap [s^-, s^+]| = 1$. Alternatively, $D(s)$ cuts e iff $|[u^-, v^-] \cap \{s^-, s^+\}| = 1$. These observations suggest that this is not a problem about graphs, but about intervals, and perhaps a problem suited for range trees.

Range trees on the Euler order. Let R be a balanced range tree with $\mathcal{V}^\pm$ at its leaves. As a balanced tree with $2n$ leaves, R has $O(n)$ nodes and height $O(\log n)$. Each range-node $a \in \mathcal{V}(R)$ induces an interval I_a on $\mathcal{V}^\pm$, consisting of all the elements of $\mathcal{V}^\pm$ at the leaves of the subtree R_a rooted at a. We call I_a a *canonical interval* of $\mathcal{V}^\pm$ (induced by a), and let $\mathcal{I} = \{I_a : a \in \mathcal{V}(R)\}$ denote the collection of all $O(n)$ canonical intervals.

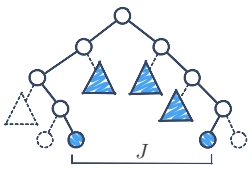

Each element in $\mathcal{V}^\pm$ appears in $O(\log n)$ canonical intervals because the height of R is $O(\log n)$. Moreover, every interval J on $\mathcal{V}^\pm$ decomposes into the disjoint union of $O(\log n)$ canonical intervals. The decomposition can be obtained in $O(\log n)$ time by tracing the two paths from the root of R to the endpoints of J and taking the canonical intervals corresponding to the $O(\log n)$ maximal subtrees between the paths (see the picture on the right). Similarly, the union of a constant number of intervals or the complement of the union of a constant number of intervals decomposes to $O(\log n)$ canonical intervals in $O(\log n)$ time.

The canonical intervals $\mathcal{I}$ relate to the 1-cuts and 2-cuts $\mathcal{C}_T$ induced by T as follows. For any pair of disjoint canonical intervals $I_1, I_2 \in \mathcal{I}$, let $\mathcal{C}(I_1, I_2) = \{(u, v) \in \mathcal{E} : u \in I_1, v \in I_2\}$ be the set of edges with one endpoint in I_1 and the other in I_2. We call $\mathcal{C}(I_1, I_2)$ a *canonical cut* induced by I_1 and I_2. We claim that any 1-cut or 2-cut $C \in \mathcal{C}_T$ decomposes into the disjoint union of $O(\log^2 n)$ canonical cuts.

For starters, let $C = \mathcal{C}(D(s))$ be the 1-cut induced by the downset of a vertex s. The 1-cut $\mathcal{C}(D(s))$ consists of all edges (u, v) with one (tagged) endpoint u^- in the interval $[s^-, s^+]$ and the other endpoint v^- outside $[s^-, s^+]$. The interval $[s^-, s^+]$ decomposes to the disjoint union $[s^-, s^+] = \bigsqcup_{I_1 \in \mathcal{I}_1} I_1$ of $O(\log n)$ canonical intervals $\mathcal{I}_1 \subseteq \mathcal{I}$, and the complement $[r^-, s^-) \cup (s^+, r^+]$ also decomposes to the disjoint union $[r^-, s^-) \cup (s^+, r^+] = \bigsqcup_{I_2 \in \mathcal{I}_2} I_2$ of $O(\log n)$ canonical intervals $\mathcal{I}_2 \subseteq \mathcal{I}$ (see the picture on the right). Together, the disjoint union $\mathcal{C}(D(s)) = \bigsqcup_{I_1 \in \mathcal{I}_1, I_2 \in \mathcal{I}_2} \mathcal{C}(I_1, I_2)$ of canonical cuts over the cross product $(I_1, I_2) \in \mathcal{I}_1 \times \mathcal{I}_2$ decomposes the 1-cut $\mathcal{C}(D(s))$ into $|\mathcal{I}_1| \times |\mathcal{I}_2| = O(\log^2 n)$ canonical cuts. Moreover, both decompositions $\mathcal{I}_1$ and $\mathcal{I}_2$ can be obtained in $O(\log n)$ time.

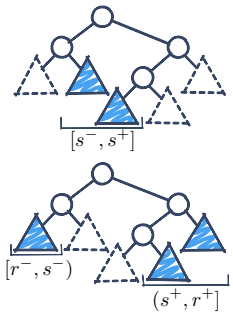

Any 2-cut decomposes similarly. If $s \parallel t$ are incomparable, with (say) $s^- < s^+ < t^- < t^+$, then the incomparable 2-cut $C = \mathcal{C}(D(s) \cup D(t))$ is the set of edges with one endpoint in $[s^-, s^+] \cup [t^-, t^+]$ and other endpoint in the complement, $[r^-, s^-) \cup (s^+, t^-) \cup (t^+, r^+]$. Both of these sets are the disjoint union of two or three intervals, and they each decompose into $O(\log n)$ canonical intervals. The cross-product of the two decompositions decomposes $\mathcal{C}(D(s) \cup D(t))$ into $O\big(\log^2 n\big)$ canonical cuts.

If $s < t$ are two comparable vertices, then $t^- < s^- < s^+ < t^+$ and the nested 2-cut $\mathcal{C}(D(t) \setminus D(s))$ is the set of edges with one endpoint in $[t^-, s^-) \cup (s^+, t^+]$ and the other endpoint in the complement $[r^-, t^-) \cup [s^-, s^+] \cup [t^+, r^+]$ and the cross-product breaks $\mathcal{C}(D(t) \setminus D(s))$ down into $O\big(\log^2 n\big)$ disjoint canonical cuts.

Thus, any 1-cut or 2-cut $C \in \mathcal{C}_T$ breaks down into to the disjoint union of $O(\log^2 n)$ canonical cuts. An edge $e = (u, v)$ appears in a canonical cut $\mathcal{C}(I_1, I_2)$ iff $u \in I_1$ and $v \in I_2$. As either end point u or v appears in $O(\log n)$ canonical intervals, e appears in at most $O(\log^2 n)$ canonical cuts. In turn, there are at most $O(m \log^2 n)$ nonempty canonical cuts. The nonempty canonical cuts are easily constructed in $O(m \log^2 n)$ time by adding each edge to the $O\big(\log^2 n\big)$ canonical cuts containing it. Taking $\mathcal{D}_T$ to be this set of nonempty canonical cuts gives Lemma 5.1.

5.2 Lazy weight increments

Lemma 5.1 identifies a relatively small set of canonical cuts $\mathcal{D}_T$ such that any 1-cut or 2-cut $C \in \mathcal{C}_T$ is the disjoint union of $O\big(\log^2 n\big)$ canonical cuts $D_1, \ldots, D_{O\left(\log^2 n\right)} \in \mathcal{D}_T$. While we can now at least gather an implicit list of all the edges of C in $O\big(\log^2 n\big)$ time, it still appears necessary to visit every edge $e \in C$ to increment its weight. The challenge is particularly tricky because the multiplicative weight updates are not uniform. However, for a fixed and static set, recent techniques by Young [2014] and Chekuri and Quanrud [2017] show how to approximately apply a multiplicative weight update to the entire set in polylogarithmic amortized time. In this section, we apply these techniques to each canonical cut $D \in \mathcal{D}_T$ and show how to combine their outputs carefully to meet the claimed bounds of Lemma 4.1.

For every canonical cut $D \in \mathcal{D}_T$, we employ the `lazy-inc` data structure of Chekuri and Quanrud [2017]. The `lazy-incs` data structure gives a clean interface to similar ideas in Young [2014], and we briefly describe a restricted version of the data structure in [Chekuri and Quanrud, 2017] per the needs of this paper[3]. The `lazy-incs` data structure is an amortized data structure that approximates a set of counters increasing concurrently at different rates. For a fixed instance of `lazy-incs` for a particular canonical cut $D \in \mathcal{D}_T$, we will have one counter for each edge tracking the "additive part" $v_e = \ln(w_e)/\epsilon$ for each edge $e \in D$. The rate of each counter e is stored as `rate`(e), and in this setting is proportional to the inverse of the capacity c_e. The primary routine is `inc`(ρ), which simulates a fractional increase on each counter proportional to a single increment at

[3]What we refer to as `lazy-incs` is most similar to `lazy-incs(0)` in Chekuri and Quanrud [2017].

the rate ρ. That is, each counter v_e increases by $\texttt{rate}(e)/\rho$. The argument ρ will always be taken to be proportional to the inverse of the minimum capacity of any edge in the cut, $\gamma = \min_{e \in C} c_e$.

The output of $\texttt{inc}(\rho)$ is a list of increments $\{(e, \delta_e)\}$ over some of the edges in D. The routine does not return an increment (e, δ_e) unless δ_e is substantially large. This allows us to charge off the work required to propagate the increment (e, δ_e) to the rest of the algorithm to the maximum weight of e. In exchange for reducing the number of increments, the sum of returned increments for a fixed edge e underestimates v_e by a small additive factor. An underestimate for $v_e = \ln(w_e)/\epsilon$ within an additive factor of $O(1)$ translates to an underestimate for $w_e = \exp(\epsilon v_e)$ within a $(1 + O(\epsilon))$-multiplicative factor, as desired.

At a high level, `lazy-incs` buckets the counters by rounding up each rate to the nearest power of 2. For each power of 2, an auxiliary counter is maintained with rate set to this power (more or less) exactly and efficiently. Counters with rates that are powers of 2 can be maintained in constant amortized time for the same reason that a binary number can be incremented in constant amortized time. When an auxiliary counter increases to the next whole integer, the tracked counters in the corresponding bucket are increased proportionately. Up to accounting details, this lazy scheme tracks the increments of each counter to within a constant additive factor at all times. A more general implementation of `lazy-incs` and a detailed proof of the following theorem are given in [Chekuri and Quanrud, 2017].

Lemma 5.2 (Chekuri and Quanrud, 2017)**.** *Let $e_1, \ldots, e_k$ be k counters with rates $\texttt{rate}(e_1) \geq \texttt{rate}(e_2) \geq \cdots \geq \texttt{rate}(e_k)$ in sorted order.*

(a) An instance of `lazy-incs` *can be initialized in $O(k)$ time.*

(b) Each `inc` *runs in $O(1)$ amortized time plus $O(1)$ for each increment returned.*

(c) `flush` *runs in $O(k)$ time.*

(d) For each counter e_i, let v_i be the true value of the counter and let $\tilde{v}_i$ be the sum of increments for counter i in the return values of `inc` *and* `flush`*.*

(i) After each call to `inc`*, we have $\tilde{v}_i \leq v_i \leq \tilde{v}_i + O(1)$ for each counter e_i.*

(ii) After each call to `flush`*, we have $\tilde{v}_i = v_i$ for each counter e_i.*

We note that `flush` is not implemented in [Chekuri and Quanrud, 2017], but is easy enough to execute by just reading off all the residual increments in the data structure and resetting these values to 0.

The application of `lazy-incs` to our problem bears some resemblance to its application to packing intervals in [Chekuri and Quanrud, 2017], where the weights are also structured by range trees. We first invoke Lemma 5.1 to generate the family of canonical cuts $\mathcal{D}_T$. For each canonical cut $D \in \mathcal{D}_T$, by Lemma 5.2, we instantiate an instance of `lazy-incs` where each edge $e \in D$ has rate $\texttt{rate}(e) = \log^2 n / c_e$. (Lemma 5.2 takes as an input a collection of counters sorted by rate. Rather than sort each canonical cut individually, we can save a $\log(n)$-factor by a slight modification to Lemma 5.1. If the edges are sorted before applying Lemma 5.1, then each canonical cut $D \in \mathcal{D}_T$ will naturally be in sorted order.) We also compute, for each canonical cut $D \in \mathcal{D}_T$, the minimum capacity $\gamma_D = \min_{e \in D} c_e$ in the cut. These preprocessing steps take $O\big(m \log^2 n\big)$ time total.

When incrementing weights along a cut $C \in \mathcal{C}_T$, we update the weights as follows. Lemma 5.1 divides C into the union of $O(\log^2 n)$ disjoint canonical cuts $\mathcal{D}'_T \subseteq \mathcal{D}_T$. We find the minimum capacity $\gamma = \min_{e \in C} c_e$ by taking the minimum of the precomputed minimum capacities of each canonical cut $D \in \mathcal{D}'_T$, i.e., $\gamma = \min_{D \in \mathcal{D}'_T} \gamma_D$. For each canonical set $D \in \mathcal{D}'_T$, we call $\texttt{inc}(\log^2 n/\gamma)$ on the corresponding `lazy-incs` instance. For every increment (e, δ) returned in `lazy-incs`, where $e \in C$ and $\delta > 0$ represents a increment to $v_e = \ln(w_e)/\epsilon$, we increase v_e by an additive factor of $\delta / \log^2 n$, and multiply w_e by $\exp\big(\epsilon\delta / \log^2 n\big)$ accordingly.

Fix an edge $e \in \mathcal{E}$. The edge e is a member of $O\big(\log^2 n\big)$ canonical cuts in $\mathcal{D}_T$, and receives its weight increments from $O(\log^2 n)$ instances of `lazy-incs`. The slight errors from each instance of `lazy-incs` accumulate additively (w/r/t $v_e = \ln(w_e)/\epsilon$). Although each instance of `lazy-incs` promises only a constant additive error in Lemma 5.2, we scaled up the rate of e from $1/c_e$ to $\log^2 n/c_e$ and divide the returned weight increments by a factor of $\log^2 n$. The scaling reduces the error from each instance of `lazy-incs` from an additive factor of $O(1)$ to $O(1/\log^2 n)$. Consequently, the sum of all $O(\log^2 n)$ instances of `lazy-incs` underestimates v_e to an $O(1)$ additive error w/r/t v_e, and underestimates w_e by at most a $(1 + O(\epsilon))$-multiplicative factor, as desired.

Increasing the sensitivity of each `lazy-incs` data structure by a factor of $\log^2 n$ means every increment (e, δ) returned by `inc` may increase v_e by as little as $1/\log^2 n$. An additive increase in $1/\log^2 n$ corresponds to a multiplicative increase of $\exp\big(\epsilon/\log^2 n\big)$ in w_e, hence property (e) in Lemma 4.1.

Together, the construction of the canonical cuts $\mathcal{D}_T$ and the $O(m \log^2 n)$ carefully calibrated instances of `lazy-inc`, one for each canonical cut $D \in \mathcal{D}_T$, give Lemma 4.1.

6 Putting it all together: Proof of Theorem 1.1

In this section we combine the ingredients discussed so far and outline the proof of Theorem 1.1. Let $(\mathcal{G}, c)$ be an Metric-TSP instance and $\epsilon' > 0$ be the error parameter. We assume without loss of generality that that ϵ' is sufficiently small, and in particular that $\epsilon' < 1/2$. We will also assume that $\epsilon' > 1/n^2$ for otherwise one could use an exact algorithm and achieve the desired time bound. This implies that $\log \frac{1}{\epsilon'} = O(\log n)$.

Consider the algorithm in Figure 5 which takes ϵ as a parameter. We will choose $\epsilon = \rho\epsilon'$ for some sufficiently small ρ. It thus suffices to argue that the algorithm outputs a $(1 + O(\epsilon))$-approximation with high probability.

High-level MWU analysis: At a high-level, our algorithm is a standard width-independent MWU algorithm for pure packing problems with an α-approximate oracle, for $\alpha = (1 + O(\epsilon))$. Our implementation follows the "time"-based algorithm of Chekuri et al. [2015]. An exact oracle in this setting repeatedly solves the global minimum cut problem w/r/t the edge weights w. To implement an α-approximate oracle, every cut output needs to be an α-approximate minimum cut.

To argue that we are indeed implementing an α-approximate minimum cut oracle, let us identify the two points at which we deviate from Karger's minimum cut algorithm [2000], which would otherwise be an exact oracle that succeeds with high probability. In both Karger's algorithm and our partially dynamic extension, we sample enough spanning trees from an approximately maximum tree packing to contain all $(1+O(\epsilon))$-approximate minimum cuts as a 1-cut or 2-cut with probability $1-\text{poly}(1/n) = 1-\text{poly}(\epsilon/n)$ (see Theorem 3.1). The sampled tree packings are the only randomized component of the algorithm, and we will argue that all the sampled tree packings succeed with high probability later after establishing a basic correctness in the event that all the sampled packings succeed.

Karger's algorithm searches all the spanning trees for the minimum 1-cut or 2-cut. We retrace the same search, except we output any approximate minimum 1-cut or 2-cut found in the search. More precisely, we maintain a target value λ with the invariant that there are no cuts of value $< \lambda$, and output any 1-cut or 2-cut with weight $\leq (1 + O(\epsilon))\lambda$ at that moment (see Lemma 4.2). Since λ is no more than the true minimum cut w/r/t w at any point, any cut with weight $\leq (1 + O(\epsilon))\lambda$ is a $(1 + O(\epsilon))$-approximation to the current minimum cut.

The second point of departure is that we do not work with the "true" weights implied by the framework, but a close approximation. By Lemma 4.1, the approximated weights underestimate

the true weights by at most a $(1 + O(\epsilon))$-multiplicative factor. In turn, we may underestimate the value of a cut by at most a $(1 + O(\epsilon))$-multiplicative factor, but crucially any 1-cut or 2-cut will still have value $\leq (1 + O(\epsilon))\lambda$, for a slightly larger constant hidden in the $O(\epsilon)$.

This firmly establishes that the proposed algorithm in fact implements an α-approximate oracle for $\alpha = 1+O(\epsilon)$ (with high probability). The MWU framework then guarantees that the final packing of cuts output by the algorithm approximately satisfies all the constraints with approximately optimal value, as desired.

Probability of failure. We now consider the probability of the algorithm failing. Randomization enters the algorithm for two reasons. Initially, we invoke Karger's randomized minimum cut algorithm [2000] to compute the value of the minimum cut w/r/t the initial weights $w = 1/c$, in order to set the first value of λ. Second, at the beginning of each epoch, we randomly sample a subset of an approximately maximum tree packing hoping that the sampled trees contain every $(1 + O(\epsilon))$-approximate minimum cut as a 1-cut or 2-cut in one of the sampled trees[4].

Karger [2000] showed that the minimum cut algorithm fails with probability at most $\text{poly}(1/n)$. He also showed (see Theorem 3.1) that each random sample of $O(\log(n/\epsilon)) = O(\log n)$ spanning trees of an approximately maximum tree packing (conducted at the beginning of an epoch) fails to capture all $(1 + \epsilon)$-approximate minimum cuts with probability at most $\text{poly}(1/n) = \text{poly}(\epsilon/n)$. As discussed in Section 3, there are a total of $O\big(\log n/\epsilon^2\big)$ epochs over the course of the algorithm. By the union bound, the probability of *any* random sample of spanning trees failing is $O\big(\log n/\epsilon^2\big) \cdot \text{poly}(\epsilon/n) = \text{poly}(\epsilon/n)$. Taking a union bound again with the chance of the initial call to Karger's minimum cut algorithm failing, the probability of any randomized component of the algorithm failing is at most $\text{poly}(1/n)$, as desired.

Running time. It remains to bound the running time of the algorithm. The analysis is not entirely straightforward, as some operations can be bounded directly, while others are charged against various upper bounds given by the MWU framework.

The algorithm initializes the edge weights in $O(m)$ time, and invokes Karger's minimum cut algorithm [2000] once which runs in $O\big(m \log^3 n\big)$ time. As established in Section 3, the remaining algorithm is divided into $O\big(\log(n)/\epsilon^2\big)$ epochs. Each epoch invokes Theorem 3.1 once to sample $O(\log(n/\epsilon)) = O(\log n)$ spanning trees from an approximate tree packing in $O\big(m \log^3 n\big)$ time. We then invoke Lemma 4.2 for each spanning tree in the sample. Over the course of $O\big(\log(n)/\epsilon^2\big)$ epochs, each with $O(\log(n))$ spanning trees, we invoke Lemma 4.2 at total of $O\big(\log^2(n)/\epsilon^2\big)$ times.

Suppose we process a tree T, outputting K approximately minimum cuts and making I full edge weight increments. By Lemma 4.2, the total time to process T is $O\big(m \log^2 n + K \log^2 n + I \log n\big)$. The first term, $O\big(m \log^2 n\big)$, comes from following the tree-processing subroutine of Karger [2000]. Over a total $O\big(\log^2(n)/\epsilon^2\big)$ trees, the total time spent mimicking Karger's subroutine is $O\big(m \log^4(n)/\epsilon^2\big)$.

The remaining terms of Lemma 4.2, depending on K and I, are introduced by the MWU framework and can be amortized against standard properties of the MWU framework. The quantity K represents the number of approximate minimum cuts output while processing a single tree T, and the sum of all K across all trees is the total number of approximate minimum cuts output. Each such cut corresponds to a single iteration, and the total number of iterations in the MWU framework is $O\big(m \log(n)/\epsilon^2\big)$. Thus, the sum of the second term $O\big(K \log^2 n\big)$ over all invocations of Lemma 4.2 is $O\big(m \log^3(n)/\epsilon^2\big)$.

The third term of Lemma 4.2, $O(I \log n)$ depends on the number of increments I returned by the `inc-cut` subroutine of the `lazy-inc-cuts` data structure. Each increment (e, δ) returned by

[4]Interestingly, the only randomization in Karger's algorithm also stems from randomly sampling an approximately maximum tree packing.

`inc-cut` increases w_e by at least a $(1+\Omega(\epsilon/\log^2 n))$-multiplicative factor (i.e., $\delta \geq \Omega(\epsilon w_e/\log^2 n)$). The MWU framework shows that a single edge can increase by a $(1+\Omega(\epsilon/\log n))$ factor at most $O\big(\log^3(n)/\epsilon^2\big)$ times (see Section 2). It follows that the total number of edge increments, over all edges, returned by `lazy-inc-cuts` is $O\big(m\log^3(n)/\epsilon^2\big)$. Thus, the sum of the quantities $O(I\log n)$ over all invocations of Lemma 4.2 is $O\big(m\log^4(n)/\epsilon^2\big)$.

All told, the algorithm runs in $O\big(m\log^4(n)/\epsilon^2\big) = \tilde{O}\big(m/\epsilon^2\big)$ time total.

Acknowledgements. We thank Neal Young for suggesting that Metric-TSP may be amenable to the techniques from [Chekuri and Quanrud, 2017].

References

P. K. Agarwal and J. Pan. Near-linear algorithms for geometric hitting sets and set covers. In *Proc. 30th Annu. Sympos. Comput. Geom.* (SoCG), page 271, 2014.

Z. Allen-Zhu and L. Orecchia. Nearly-linear time positive LP solver with faster convergence rate. In *Proc. 47th Annu. ACM Sympos. Theory Comput.* (STOC), pages 229–236, 2015.

D. L. Applegate, R. E. Bixby, V. Chvatal, and W. J. Cook. *The traveling salesman problem: a computational study.* Princeton University Press, 2011.

S. Arora, E. Hazan, and S. Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing*, 8(1):121–164, 2012.

Y. Bartal and L.-A. Gottlieb. A linear time approximation scheme for euclidean tsp. In *Proc. 54th Annu. IEEE Sympos. Found. Comput. Sci.* (FOCS), pages 698–706, 2013.

C. Chekuri and K. Quanrud. Near-linear time approximation schemes for some implicit fractional packing problems. To appear in Proc. 28th ACM-SIAM Sympos. Discrete Algs. *(SODA)*, 2017, 2017.

C. Chekuri, T. Jayram, and J. Vondrák. On multiplicative weight updates for concave and submodular function maximization. In *Proc. 6th Conf. Innov. Theoret. Comp. Sci.* (ITCS), pages 201–210, 2015.

N. Christofides. Worst-case analysis of a new heuristic for the traveling salesman problem. Technical Report 388, Graduate School of Industrial Administration, Carnegie Mellon University, 1976.

W. J. Cook. *In purusit of the traveling salesman: mathematics at the limits of computation.* Princeton University Press, 2014.

G. B. Dantzig, D. R. Fulkerson, and S. M. Johnson. Solution of a large-scale traveling-salesman problem. *Operations Research*, 2(4):393–410, 1954.

L. K. Fleischer. Approximating fractional multicommodity flow independent of the number of commodities. *SIAM J. Discrete Math.*, 13(4):505–520, 2000. Preliminary version in Proc. 40th Annu. IEEE Sympos. Found. Comput. Sci. *(FOCS)*, 1999.

H. N. Gabow. A matroid approach to finding edge connectivity and packing arborescences. *J. Comput. Sys. Sci.*, 50(2):259–273, 1995. Preliminary version in Proc. 23rd Annu. ACM Sympos. Theory Comput. *(STOC)*, 1991.

N. Garg and J. Könemann. Faster and simpler algorithms for multicommodity flow and other fractional packing problems. *SIAM J. Comput.*, 37(2):630–652, 2007. Preliminary version in Proc. 39th Annu. IEEE Sympos. Found. Comput. Sci. *(FOCS)*, 1998.

M. X. Goemans. Worst-case comparison of valid inequalities for the TSP. *Math. Prog.*, 69(1-3): 335–349, 1995.

M. X. Goemans and D. Bertsimas. Survivable networks, linear programming relaxations and the parsimonious property. *Math. Prog.*, 60(1):145–166, June 1993.

M. D. Grigoriadis and L. G. Khachiyan. Fast approximation schemes for convex programs with many blocks and coupling constraints. *SIAM J. Optim.*, 4(1):86–107, 1994.

G. Gutin and A. P. Punnen. *The traveling salesman problem and its variations*, volume 12. Springer Science & Business Media, 2006.

S. Har-Peled. *Geometric Approximation Algorithms*, volume 173 of *Mathematical Surveys and Monographs*. Amer. Math. Soc., 2011.

D. Harel and R. E. Tarjan. Fast algorithms for finding nearest common ancestors. *SIAM J. Comput.*, 13(2):338–355, 1984.

M. Held and R. M. Karp. The traveling-salesman problem and minimum spanning trees. *Operations Research*, 18(6):1138–1162, 1970.

M. R. Henzinger and V. King. Randomized fully dynamic graph algorithms with polylogarithmic time per operation. *J. Assoc. Comput. Mach.*, 46(4):502–516, 1999. Preliminary version in Proc. 27th Annu. ACM Sympos. Theory Comput. *(STOC)*, 1995.

D. R. Karger. Random sampling and greedy sparsification for matroid optimization problems. *Math. Program.*, 82:41–81, 1998. Preliminary version in Proc. 34th Annu. IEEE Sympos. Found. Comput. Sci. *(FOCS)*, 1993.

D. R. Karger. Minimum cuts in near-linear time. *J. Assoc. Comput. Mach.*, 47(1):46–76, 2000. Preliminary version in Proc. 28rd Annu. ACM Sympos. Theory Comput. *(STOC)*, 1996.

R. Khandekar. *Lagrangian relaxation based algorithms for convex programming problems.* PhD thesis, Indian Institute of Technology Delhi, Mar. 2004.

C. Koufogiannakis and N. E. Young. A nearly linear-time PTAS for explicit fractional packing and covering linear programs. *Algorithmica*, 70(4):648–674, 2014. Preliminary version in Proc. 48th Annu. IEEE Sympos. Found. Comput. Sci. *(FOCS)*, 2007.

M. Lampis. Improved inapproximability for TSP. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 243–253. Springer, 2012.

E. L. Lawler, J. K. Lenstra, A.-G. Rinnooy-Kan, and D. B. Shmoys. *The traveling salesman problem.* John Wiley & Sons, Ltd., 1985.

C. L. Monma, B. S. Munson, and W. R. Pulleyblank. Minimum-weight two-connected spanning networks. *Math. Prog.*, 46(1):153–171, Jan. 1990.

A. Mądry. Faster approximation schemes for fractional multicommodity flow problems via dynamic graph algorithms. In *Proc. 42nd Annu. ACM Sympos. Theory Comput.* (STOC), pages 121–130, 2010.

C. S. J. A. Nash-Williams. Edge-disjoint spanning trees of finite graphs. *J. London Math. Soc.*, 36: 445–450, 1961.

R. Peng. Approximate undirected maximum flows in $o(mpolylog(n))$ time. In *Proc. 27th ACM-SIAM Sympos. Discrete Algs.* (SODA), pages 1862–1867, 2016.

S. A. Plotkin, D. B. Shmoys, and É. Tardos. Fast approximation algorithms for fractional packing and covering problems. *Math. of Oper. Res.*, 20(2):257–301, 1995.

S. B. Rao and W. D. Smith. Approximating geometrical graphs via "spanners" and "banyans". In *Proc. 30th Annu. ACM Sympos. Theory Comput.* (STOC), pages 540–550, 1998.

D. Sleator and R. E. Tarjan. A data structure for dynamic trees. *J. Comput. Syst. Sci.*, 26:362–391, 1983.

R. E. Tarjan and U. Vishkin. An efficient parallel biconnectivity algorithm. *SIAM J. Comput.*, 14 (4):862–874, 1985.

M. Thorup. Fully-dynamic min-cut. *Combinatorica*, 27(1):91–127, 2007. Preliminary version in Proc. 33rd Annu. ACM Sympos. Theory Comput. *(STOC)*, 2001.

W. T. Tutte. On the problem of decomposing a graph into n connected components. *J. London Math. Soc.*, 36:221–230, 1961.

J. Vygen. New approximation algorithms for the TSP. *OPTIMA*, 90:1–12, 2012.

D. Wang, S. Rao, and M. Mahoney. Unified acceleration method for packing and covering problems via diameter reduction. *CoRR*, abs/1508.02439, August 2015. To appear in Proc. 43rd Internat. Colloq. Automata Lang. Prog. *(ICALP)*, 2016.

N. E. Young. Nearly linear-time approximation schemes for mixed packing/covering and facility-location linear programs. *CoRR*, abs/1407.3015, 2014. URL `http://arxiv.org/abs/1407.3015`.

A Computing the primal solution

Our MWU algorithm computes a $(1-\epsilon)$-approximate solution to 2ECSSD. Recall that the algorithm maintains a weight $w(e)$ for each edge e. These weights evolve with time. Using a standard argument we show how we can recover a $(1+\epsilon)$-approximation to the LP 2ECSS.

Let OPT be the optimum solution value of 2ECSSD and 2ECSS. Our algorithm computes an implicit packing in the LP 2ECSSD of value $(1-\epsilon)\,\mathrm{OPT}$ in near-linear time. It does this via the MWU method which works in several iterations. In each iteration i the algorithm computes a global mincut C_i with respects to the edge weights $w_i(e)$ where $w_i(e)$ is the weight of edge e in iteration i. Each iteration corresponds to solving a relaxation of the original LP by collapsing the constraints into a single constraint via the weights in iteration i. The value of the solution to this relaxation in iteration i is $(\sum_e w_i(e)c(e))/w_i(C_i)$. Since this is a relaxation to the original LP, the value of this solution is at least OPT. Over the entire algorithm we have a time of one unit, and the integral of

all solutions we compute is at least OPT. Let z be the averaged solution at the end of the algorithm. Via the MWU analysis z violates each constraint by at most a $(1+\epsilon)$-factor. Thus by scaling down z by $(1+\epsilon)$, we ensure feasibility of constraints and obtain a feasible solution to 2ECSSD of value at least $\frac{1}{1+\epsilon}\,\mathrm{OPT}$.

We claim (and this is standard) that there is an iteration i such that the value of the solution $\sum_e w_i(e)c(e)/w_i(C_i) \leq (1+\epsilon)\,\mathrm{OPT}$. If not, the value of z would be larger than $(1+\epsilon)\,\mathrm{OPT}$ and we would, after scaling down by a factor of $(1+\epsilon)$, obtain a feasible solution to 2ECSSD of value more than OPT, a contradiction. Given an iteration i where $\sum_e w_i(e)c(e)/w_i(C_i) \leq (1+\epsilon)\,\mathrm{OPT}$ we can obtain a feasible solution to 2ECSS as follows. Simply set $y(e) = w_i(e)/w_i(C_i)$ for each e; we claim that this is a feasible solution whose cost is at most $(1+\epsilon)\,\mathrm{OPT}$. To see that y is feasible, consider any cut C_j. The capacity of the cut is $\sum_{e \in C_j} y_e = w_i(C_j)/w_i(C_i) \geq 1$ since C_i is the cheapest cut with respect to the weights $w_i(e)$. The cost of the solution y is $\sum_e y(e)c(e) = \sum_e w_i(e)c(e)/w_i(C_i) \leq (1+\epsilon)\,\mathrm{OPT}$ by choice of i.

The solution y can be output in nearly-linear time. Our algorithm maintains the weights $w_i(e)$ implicitly via data structures. However, during iteration i we know the value of the mincut $w_i(C_i)$ to within a $(1+O(\epsilon))$ factor as well as the sum $\sum_e w_i(e)c(e)$. In particular, we can keep track of the iteration i^* that minimizes this ratio; this iteration i^* will satisfy the desired condition, $\sum_e w_{i^*}(e)c(e)/w_{i^*}(C_{i^*}) \leq (1+O(\epsilon))\,\mathrm{OPT}$. The actual set of weights $w_{i^*}(e)$ needed to create the solution $y(e)$ can be recovered by either replaying the algorithm (by storing the necessary randomization from the first run) or simply maintaining the weights by persistent data structures and keeping a pointer to the persistent copy of the weights from iteration i^*.